Quantum entanglement at the energy frontier: Measurement of spin entanglement in top-antitop quark pairs with the ATLAS experiment

Dissertation
zur Erlangung des Doktorgrades
an der Fakultät für Mathematik, Informatik und Naturwissenschaften
Fachbereich Physik
der Universität Hamburg

vorgelegt von
Fiona Ann Jolly

Hamburg
2026

Gutachter/innen der Dissertation:	Prof. Dr. Peter Schleper Dr. Janna Katharina Behr
Zusammensetzung der Prüfungskommission:	Prof. Dr. Peter Schleper Dr. Janna Katharina Behr Dr. Jürgen Reuter Prof. Dr. Elisabetta Gallo Prof. Dr. Arwen Pearson
Vorsitzende/r der Prüfungskommission:	Prof. Dr. Arwen Pearson
Datum der Disputation:	13.08.2026
Vorsitzender des Fach-Promotionsausschusses PHYSIK:	Prof. Dr. Johannes Haller
Leiter des Fachbereichs PHYSIK:	Prof. Dr. Markus Drescher
Dekan der Fakultät MIN:	Prof. Dr.-Ing. Norbert Ritter

Declaration of Oath

I hereby declare and affirm that this doctoral dissertation is my own work and that I have not used any aids and sources other than those indicated.

If electronic resources based on generative artificial intelligence (gAI) were used in the course of writing this dissertation, I confirm that my own work was the main and value-adding contribution and that complete documentation of all resources used is available in accordance with good scientific practice. I am responsible for any erroneous or distorted content, incorrect references, violations of data protection and copyright law or plagiarism that may have been generated by the gAI.

01.06.2026	
Date	Signature of doctoral candidate

Acknowledgements

As this thesis marks the completion of my doctoral journey, I would like to reflect on the people and experiences that have made it possible. The list of people I would like to thank would go much longer than these pages allow, so not everyone is mentioned. Nevertheless, I am deeply grateful to each and every one of you.

First and foremost, I would like to express my sincere gratitude to my supervisor, Katharina Behr, for being a true inspiration to me. The enthusiasm, dedication, and passion with which you approach your work have always awed me. It goes without saying that a significant part of my PhD journey has been moulded by your exceptional supervision. Thank you for always being approachable, for giving me the freedom to think independently and present my own ideas, and for consistently encouraging me to achieve my best. Your ability to juggle countless responsibilities while remaining calm and composed is truly remarkable, and it is a quality that has certainly rubbed off on me to some extent. I am also deeply grateful for the many opportunities you have given me to participate in schools, workshops, and conferences, all of which have greatly enriched my PhD experience. *Thank you!*

Thanks to Peter Schleper for reviewing my thesis and also for the fruitful discussions during the yearly report meetings. Thank you to Jürgen Reuter, Elisabetta Gallo, and Arwen Pearson for agreeing to be members of my doctoral committee.

I would like to thank Nicola de Biase and Marcus Vinicius Gonzalez Rodrigues, who were there during the initial stages of my PhD journey, for taking the time to help me with my analysis and qualification project.

A huge thank you goes to my postdoctoral supervisor, Eleanor Jones. Thank you for always being available to answer my questions whenever I knock on your door, even at the last minute. Beyond your contributions to this work, I am thankful for your guidance, encouragement, and support on countless occasions over the years, both professionally and personally.

I would also like to thank Stephen Jiggins for the valuable discussions that helped me shape the statistical analysis framework of the boosted analysis. Apart from that, the enjoyable conversations and stories we shared outside of work, made the journey all the more pleasant!

The work presented in this thesis has been carried out in collaboration with many researchers, all of whom have contributed in different ways. I would like to thank Theresa Reisch for helping me become familiar with the threshold entanglement analysis framework and for being a pleasure to work with. I am also grateful to Andrea Knue for providing essential inputs at various stages of my analysis. My thanks also go to Baptiste Ravina, Yoav Afik, and Federica Fabbri for their valuable feedback and constructive discussions throughout the course of this work. Furthermore, I would like to thank Sylvia Louise Mason and Nicola Forti for their support in providing the essential inputs for the statistical interpretation of the boosted analysis. A special thanks goes to Ivo Young for being not only a great

analysis partner but also a good friend.

The international environment at DESY, shaped by countless colleagues, staff members, and researchers, played an important role in making my research experience both productive and enjoyable. I would like to thank Ingrid Gregor, the ATLAS group leader, for her support with various administrative and official matters throughout my time at DESY. I am also grateful to the secretaries, Sabine Krohn and Carolin Groetzinger, who always seemed to have a quick answer to every organisational or paperwork-related question I had.

I feel truly fortunate to have made so many friends at work who have supported me in some way or another. Sara, we both have come a long way together since we began our Masters at Bonn. Thanks for being such a wonderful human being. Your knocks on my office door and calls for coffee breaks often provided a welcome respite during some of the busiest days at work. And I still believe that I am taller than you! Dario, your regular visits to my office, often accompanied by chocolates, will always be remembered. Thank you for being a patient listener and a great person to have around. Max, thank you for being that carefree and unfiltered friend with whom I could always speak openly, whether about work or just the random things in life. Manuel, you are one of the sweetest people I know, and somehow we seem to share the same taste buds when it comes to food! Thank you for all the "You've got this" encouragement throughout the years. Serhat, thank you for all the support you have given me throughout different stages of my PhD, and for being such a wonderful neighbour with whom I had the pleasure of sharing a few baking adventures. Supriya, I cannot thank you enough for all your help and support, both within and beyond work. I feel fortunate to have known you, even before I started my PhD. Gediminas ("G"), I will always remember you as the one who helped me become a tour guide. Thank you for also being available to help whenever needed. The board game nights with you and many others were always so much fun. Thanks to Weronika, Rebecca, Bhupesh, Björk, Siemen, and Henry for making the office a pleasant place to work. Thanks Daniel for kindly taking the time to read a part of my thesis and for the helpful feedback. To each and every member of the ATLAS DESY group *Thank you!*. And to all my friends at badminton, another big thank you for creating an atmosphere where I could forget about everything else for a while and simply enjoy the game.

Before coming to DESY, I was fortunate to make many friends during my Master's studies in Bonn who have remained in touch with me and continuously motivated me throughout my PhD. I would like to thank Murillo, Fadi, Vlad, Celia, Gulnaz, Rishabh, Aditya, Anahat, Aakash, Anushka, Rajat, and especially Anubhav for always looking out for me. Taking time to visit some of you in Bonn during my PhD years always brought a smile to my face and reminded me of many fond memories. Continuing on that note of nostalgia, I would also like to thank my friends from college in India, who have likewise supported me and stayed in touch over the years. Thank you, Sharlet, Maria, Eby, Mathew, Mervin, Mariya, and Deepthi, for your friendship and encouragement.

I would like to thank all my relatives, especially Renny Uncle, for always believing in me and for encouraging me to achieve more than I thought possible. A big thank you to Elke, who was a tremendous help in settling me into Hamburg during the early stages of my PhD.

Finally, my deepest gratitude goes to those to whom I dedicate this thesis: *my family*. They are the reason I have been able to achieve so much. Even though my parents do not really understand the language of my work, they have never faltered in expressing the support and motivation that has kept me going. Thank you, Appa and Amma. Thanks to my sister and best friend, Ciona, who, just like our parents, has been my forever support system.

To those reading this thesis for the first time, I would like to borrow a phrase from The Book Thief and say: "I have hated the words and I have loved them, and I hope I have made them right."

Abstract

In this thesis, measurements of quantum entanglement between the spins of pair-produced top and antitop quarks ($t\bar{t}$) are presented. The sensitivity of these measurements is found to be highest in two distinct kinematic regions. The first region is characterised by low values of the invariant mass of the $t\bar{t}$ system, $m_{t\bar{t}}$, close to the production threshold, where $t\bar{t}$ pairs are expected to be predominantly produced in a maximally entangled spin-singlet state. The second region corresponds to very high $m_{t\bar{t}}$, where the top quarks are highly boosted and are predicted to form a maximally entangled spin-triplet state. Accordingly, two complementary analyses are performed, targeting entanglement measurements in these distinct regions of phase space.

The measurements are performed using proton-proton collision data collected by the ATLAS experiment at the Large Hadron Collider at CERN. The dataset corresponds to an integrated luminosity of 140 fb^{-1}, recorded during the 2015-2018 data-taking period at a centre-of-mass energy of $\sqrt{s}$ = 13 TeV.

The analyses focus on $t\bar{t} \rightarrow (W^{+}b)(W^{-}\bar{b})$ final states in which one W boson decays leptonically to an electron or muon and the corresponding (anti)neutrino, and the other decays hadronically to a quark-antiquark pair. These events are characterised experimentally by the presence of a charged lepton, multiple hadronic jets, and missing transverse momentum. This topology is referred to as the lepton+jets channel.

Under the assumption of the Standard Model (SM), the entanglement measurement near the $t\bar{t}$ production threshold is performed using an observable D, inferred from an appropriately defined angle between the decay products of the two top quarks that retain maximal spin information. The data are compared to SM predictions, including both standard $t\bar{t}$ production and, for the first time in an ATLAS entanglement measurement, contributions from colour-singlet quasi-bound-state formation near threshold. Entanglement is observed with a significance of 3.73(3.97) standard deviations with respect to the hypothesis of separable (i.e non-entangled) top quarks.

At high $m_{t\bar{t}}$, entanglement is probed using a related observable, D_3, defined analogously from an angular variable. This study constitutes the first ATLAS entanglement measurement specifically targeting the high-$m_{t\bar{t}}$ regime and is part of an ongoing ATLAS analysis effort. Entanglement is also expected to be observed in this kinematic region, providing an important extension of entanglement measurements into a regime not yet explored by quantum information measurements in ATLAS.

Zusammenfassung

In dieser Dissertation werden Messungen der Quantenverschränkung zwischen den Spins von paarweise erzeugten Top- und Antitop-Quarks ($t\bar{t}$) vorgestellt. Es zeigt sich, dass die Empfindlichkeit dieser Messungen in zwei unterschiedlichen kinematischen Regionen am höchsten ist. Die erste Region ist durch niedrige Werte der invarianten Masse des $t\bar{t}$-Systems ($m_{t\bar{t}}$) nahe der Produktionsschwelle gekennzeichnet; hier wird erwartet, dass $t\bar{t}$-Paare vorwiegend in einem maximal verschränkten Spin-Singulett-Zustand erzeugt werden. Die zweite Region entspricht sehr hohen Werten von $m_{t\bar{t}}$, in denen die Top-Quarks stark geboostet sind und - der Vorhersage zufolge - einen maximal verschränkten Spin-Triplett-Zustand bilden. Dementsprechend werden zwei komplementäre Analysen durchgeführt, die auf Verschränkungsmessungen in diesen spezifischen Bereichen des Phasenraums abzielen.

Die Messungen basieren auf Daten aus Proton-Proton-Kollisionen, die vom ATLAS-Experiment am Large Hadron Collider (LHC) am CERN aufgezeichnet wurden. Der Datensatz entspricht einer integrierten Luminosität von $140\,\text{fb}^{-1}$ und wurde während der Datennahmeperiode 2015-2018 bei einer Schwerpunktsenergie von $\sqrt{s} = 13$ TeV erfasst.

Die Analysen konzentrieren sich auf $t\bar{t} \to (W^+b)(W^-\bar{b})$-Endzustände, bei denen ein W-Boson leptonisch in ein Elektron oder Myon sowie das entsprechende (Anti-)Neutrino zerfällt, während das andere hadronisch in ein Quark-Antiquark-Paar zerfällt. Experimentell sind diese Ereignisse durch das Auftreten eines geladenen Leptons, mehrerer hadronischer Jets sowie fehlenden Transversalimpulses gekennzeichnet. Diese Topologie wird als „Lepton+Jets“-Kanal bezeichnet.

Unter der Annahme des Standardmodells (SM) wird die Verschränkungsmessung nahe der $t\bar{t}$-Produktionsschwelle mithilfe einer Observablen D durchgeführt. Diese wird aus einem entsprechend definierten Winkel zwischen den Zerfallsprodukten der beiden Top-Quarks abgeleitet, welche die maximale Spin-Information bewahren. Die Daten werden mit den SM-Vorhersagen verglichen; diese umfassen sowohl die standardmäßige $t\bar{t}$-Produktion als auch - erstmals im Rahmen einer ATLAS-Verschränkungsmessung - Beiträge aus der Bildung farbsingulettartiger, quasi-gebundener Zustände nahe der Produktionsschwelle. Verschränkung wird mit einer Signifikanz von 3,73 (3,97) Standardabweichungen gegenüber der Hypothese separabler (d. h. nicht verschränkter) Top-Quarks beobachtet.

Im Bereich hoher $m_{t\bar{t}}$-Werte wird die Verschränkung mithilfe einer verwandten Observablen, D_3, untersucht, die analog auf der Grundlage einer Winkelvariablen definiert ist. Diese Studie stellt die erste ATLAS-Messung zur Verschränkung dar, die gezielt auf den Bereich hoher $m_{t\bar{t}}$-Werte ausgerichtet ist, und ist Teil einer laufenden ATLAS-Analyse. Es wird erwartet, dass Verschränkung auch in diesem kinematischen Bereich beobachtet werden kann; dies stellt eine bedeutende Erweiterung der Verschränkungsmessungen in einen Bereich dar, der im Rahmen von quanteninformationstheoretischen Messungen bei ATLAS bislang noch nicht erschlossen wurde.

Contents

1 Introduction **1**

2 Fundamentals of particle physics **7**
2.1 The Standard Model of particle physics . . . 7
2.1.1 Fundamental particles . . . 8
2.1.2 Fundamental interactions . . . 9
2.1.3 The Higgs mechanism . . . 15
2.2 Limitations of the Standard Model . . . 19

3 Top quarks: probing the intersection of particle physics and quantum theory **21**
3.1 Production and decay mechanisms of top quarks in the SM . . . 22
3.1.1 Top-quark pair production . . . 22
3.1.2 Top-quark decays . . . 25
3.2 Quantum properties of top quarks . . . 26
3.2.1 Top-quark spin analysers . . . 28
3.2.2 Spin correlations in top-quark pairs . . . 33
3.2.3 Quantum entanglement in top-quark pairs . . . 34
3.2.4 Measuring Bell correlated states with top quarks . . . 38
3.3 Exploring quasi-bound-state effects in the threshold region . . . 40

4 The ATLAS experiment at the LHC **45**
4.1 The Large Hadron Collider . . . 46
4.1.1 Essential collider parameters . . . 47
4.2 The ATLAS experiment . . . 50
4.2.1 Inner Detector . . . 52
4.2.2 Calorimeters . . . 54
4.2.3 Magnet system . . . 56
4.2.4 Muon Spectrometer . . . 57
4.2.5 Trigger and data aquisition system . . . 58
4.2.6 Data collection and analysis . . . 59
4.3 Future prospects and upgrades . . . 59

5 Data and simulation of proton-proton collisions in ATLAS **61**
5.1 ATLAS dataset . . . 62
5.2 Simulation of events . . . 62

5.3 Levels of event description . . . 64
5.3.1 Detector level . . . 64
5.3.2 Parton level . . . 65
5.3.3 Particle level . . . 65
5.4 Simulated samples . . . 65
5.4.1 SM $t\bar{t}$ production . . . 67
5.4.2 Single-top production . . . 67
5.4.3 W+jets and Z+jets . . . 68
5.4.4 $t\bar{t}V$ and $t\bar{t}H$ . . . 69
5.4.5 Diboson . . . 69
5.4.6 $t\bar{t}$ quasi-bound states . . . 69

6 Reconstruction and identification of physics objects in the ATLAS detector 71
6.1 Charged-particle track reconstruction . . . 72
6.2 Electrons . . . 78
6.3 Muons . . . 81
6.4 Jets . . . 83
6.4.1 Input objects for reconstruction . . . 83
6.4.2 Jet clustering algorithms . . . 84
6.4.3 Small-R jets . . . 86
6.4.4 Reclustered jets . . . 86
6.4.5 Jet calibration . . . 87
6.4.6 Flavour tagging . . . 89
6.5 Missing transverse momentum . . . 93
6.6 Overlap removal . . . 95

7 Probing entanglement close to the $m_{t\bar{t}}$ threshold 97
7.1 Analysis strategy . . . 97
7.2 Baseline object selection . . . 98
7.2.1 Leptons . . . 98
7.2.2 Small-radius jets . . . 98
7.3 Event selection . . . 99
7.3.1 Preselection . . . 99
7.3.2 Trigger selection . . . 100
7.3.3 Common object-related selection . . . 100
7.4 Reconstruction of the $t\bar{t}$ system . . . 101
7.4.1 SPANet . . . 102
7.4.2 Kinematic reconstruction of the $t\bar{t}$ system . . . 105
7.5 Particle-level selection and reconstruction of the $t\bar{t}$ system . . . 106
7.6 Signal and validation region definitions . . . 108
7.6.1 Reconstruction level . . . 108
7.6.2 Particle level . . . 111
7.6.3 Parton level . . . 111
7.7 Calibration curve approach . . . 111
7.8 Non-prompt and fake lepton background estimation . . . 115

7.9 Systematic uncertainties . . . 118
7.9.1 Experimental . . . 118
7.9.2 pQCD $t\bar{t}$ modelling systematics . . . 121
7.9.3 Modelling systematics of $t\bar{t}$ quasi-bound-state effects . . . 124
7.9.4 Background systematics . . . 125
7.10 Distributions and event yield . . . 126
7.11 Results . . . 129
7.11.1 Calibration curves and extraction of D using the nominal signal model . . . 129
7.11.2 Results without quasi-bound-state effects . . . 133
7.11.3 Sensitivity to entanglement . . . 135
7.11.4 Results using the simplified model of quasi-bound-state effects . . . 137
7.11.5 Comparison with other LHC entanglement measurements . . . 139

8 Probing entanglement with boosted top quarks 143
8.1 Analysis strategy . . . 144
8.2 Baseline object selection . . . 144
8.2.1 Leptons . . . 144
8.2.2 Jets . . . 145
8.3 Event selection . . . 145
8.4 Event reconstruction . . . 146
8.5 Signal regions definition . . . 148
8.6 Profile-likelihood unfolding . . . 149
8.6.1 Profile-likelihood fits . . . 149
8.6.2 Unfolding . . . 153
8.6.3 Reparametrisation and extraction of D_3 . . . 157
8.7 Systematic uncertainties . . . 157
8.7.1 Experimental . . . 158
8.7.2 $t\bar{t}$ modelling . . . 159
8.8 Event yields and differential distributions in the signal regions . . . 159
8.9 Results . . . 163
8.9.1 Expected value of D_3 . . . 163
8.9.2 Impact of systematic uncertainties in the fit . . . 166
8.10 Comparison with entanglement measurement by CMS . . . 176
8.11 Outlook . . . 177

9 Conclusion and outlook 179

A Kinematic fits for the inputs to the particle-level chi2 reconstruction 183

B Optimisation studies for entanglement signal region 185

C Unfolding inputs in muon region 187

Bibliography 191

List of Figures 211

List of Tables 219

CHAPTER 1

Introduction

"Science progresses best when observations force us to alter our preconceptions."

— Vera Rubin

The modern world of the twenty-first century is the result of a long sequence of developments that began with the earliest attempts of humankind to understand nature. Driven by curiosity and desire to comprehend the physical world, profound advances have been achieved over the past centuries across a wide range of fields, including medicine, industry, technology, and science. Among these disciplines, physics is an important branch of science which plays a fundamental role by addressing some of the most basic questions about reality: where we come from, what the universe is made of, and how it evolves. Appropriately, the word physics derives from the Latin word *physica*, meaning *the study of nature*. What began in ancient Greece during the sixth and fifth centuries BC, when early philosophers first sought rational explanations for the natural world, evolved significantly through the work of thinkers such as Francis Bacon, who introduced the concept of controlled experimentation over purely abstract reasoning based on observation alone. Later, works such as the study of motion by Isaac Newton, famously inspired by a falling apple in the seventeenth century, helped establish physics as a distinct scientific discipline, one that has since grown into a vast and ever-expanding body of knowledge, which continues to shape how we understand and interact with the world.

Pioneering work by scientists such as Galileo Galilei, Johannes Kepler, and Isaac Newton in the sixteenth and seventeenth centuries established a coherent framework for describing motion, forces, and gravitation. With the advent of the Industrial Revolution in the nineteenth century, the growing demand for energy and technological development marked the beginning of research efforts in thermodynamics, chemistry, optics, and electromagnetism. Together, these developments formed what is now referred to as ***classical physics***. While remarkably successful in describing macroscopic systems involving objects moving at relatively low speeds, classical theories proved inadequate for explaining phenomena occurring at atomic and subatomic length scales or at velocities approaching the speed of light.

The turn of the twentieth century marked the origin of ***modern physics***, driven by the foundational work of Max Planck in quantum theory and Albert Einstein in the theory of relativity. ***Quantum mechanics (QM)***, pioneered by Werner Heisenberg, Erwin Schrödinger, and Paul Dirac, provided a

successful framework for describing physical phenomena at atomic and subatomic scales. In parallel, ***special relativity*** revolutionised the understanding of space, time, and energy, becoming essential for the description of systems at very high energies.

One of the major breakthroughs in understanding the subatomic world came with the discovery of the electron by J.J. Thomson, followed by the identification of the nucleus by Ernest Rutherford [1] and the subsequent discovery of the proton and neutron [2]. These findings established the atomic model of matter, described by a small, dense nucleus surrounded by electrons. Initially, these particles were regarded as the elementary building blocks of nature, leading to the emergence of ***particle physics***. However, advancements in experimental techniques hinted that protons and neutrons were not fundamental after all. This was confirmed in the late 1960s through deep inelastic scattering experiments at the Stanford Linear Accelerator Center (SLAC), where high-energy electrons were directed at protons [3, 4]. The scattering patterns clearly revealed that protons contained point-like constituents, which are now called quarks, providing the first experimental evidence of their existence. Alongside quarks, other classes of elementary particles, such as leptons, including the electrons and neutrinos, were identified.

This growing reliance on high-energy accelerators to explore smaller length scales led to the development of ***high-energy physics (HEP)*** as the dominant experimental framework of particle physics. By the mid-twentieth century, particle physics and high-energy physics had become tightly intertwined, providing a unified approach for studying the fundamental constituents of matter and their interactions. This framework is encapsulated in the ***Standard Model (SM)*** of particle physics, developed in the 1970s. It provides a comprehensive description of all known elementary particles and their interactions through three of the four fundamental forces of nature, namely the strong, weak, and electromagnetic interactions.

Most of the predicted components of the model have been experimentally confirmed at high-energy particle accelerators. Notable examples include the discovery of the gluon, the mediator of the strong interactions, at the PETRA collider at the Deutsches Elektronen-Synchrotron (DESY) in 1979 [5], the observation of the W and Z bosons responsible for the weak interaction, at the CERN Super Proton Synchrotron (SPS) in 1983 [6, 7], and the discovery of the top quark, the heaviest known elementary particle, at the Tevatron collider in Fermilab in 1995 [8, 9]. A central feature of the SM is the Higgs field, a scalar quantum field that permeates the universe and gives mass to elementary particles through the mechanism of electroweak symmetry breaking. The discovery of the Higgs boson, the quantum excitation of this field, by the ATLAS and CMS experiments at the ***Large Hadron Collider (LHC)*** in 2012 [10, 11], marked a milestone in particle physics. Despite its remarkable success, the SM is known to be incomplete, as it does not incorporate gravity and fails to explain phenomena such as dark matter, among other open questions. Addressing these limitations requires the development of theories beyond the SM (BSM), together with new experimental studies and continued advances in detector technology to probe previously unexplored energy scales and interactions.

Among the particles in the SM, the ***top quark*** occupies a unique position. Owing to its large mass, it couples most strongly to the Higgs boson and is therefore expected to play a key role in theories of electroweak symmetry breaking and potential extensions of the SM. Moreoever, unlike all other quarks, the top quark decays before it can hadronise, a process in which quarks bind together to form composite particles known as hadrons. This unique property of the top quark allows direct access to its spin state, an intrinsic quantum-mechanical form of angular momentum, through the angular distributions of its decay products.

Owing to this unique feature, one of the most important aspects of top-quark physics is the study of

spin correlations in pair-produced top and antitop quarks ($t\bar{t}$). In recent years, ***quantum entanglement*** and other quantum effects in $t\bar{t}$ systems have also attracted significant attention. Quantum entanglement is a distinctive feature of QM, describing correlations between two particles such that the state of one cannot be fully determined without knowledge of the state of the other. Once dismissed by Einstein as "spooky action at a distance", entanglement is now recognised as a foundational resource in quantum information science. A more stringent test of QM is provided by the Bell inequality, formulated by John Bell in 1964 [12], in the context of the earlier Einstein-Podolsky-Rosen (EPR) paradox [13]. The EPR paradox questioned the completeness of QM, highlighting strong correlations between spatially separated particles that could not be explained by any classical theory based on hidden variables acting only within the immediate surroundings of each particle, known as local hidden-variable theory. The Bell inequality provides a rigorous test to distinguish such strong quantum correlations from those allowed in any local hidden-variable theory. The significance of these ideas was highlighted by the 2022 Nobel Prize in Physics, awarded to Alain Aspect, John Clauser, and Anton Zeilinger for their groundbreaking experiments with entangled photons [14], which confirmed violation of Bell inequalities and provided direct evidence of non-classical correlations. This demonstrated that QM is both complete and non-local. Historically, experimental tests of QM have been carried out in low-energy systems using photons, ions, or electrons [14–16]. With the high energies and capabilities of the LHC, it is now possible for the first time, to study such quantum effects using fundamental quarks at TeV-scale energies.

At the LHC, top quarks are predominantly produced in pairs, providing an ideal environment to probe quantum entanglement through angular observables constructed from the top-quark decay products, which are sensitive to the respective top-quark spin states. The $t\bar{t}$ production mechanisms at the LHC give rise to distinct $t\bar{t}$ spin configurations and enhance sensitivity to entangled spin states, particularly near the $t\bar{t}$ production threshold where the invariant mass of the $t\bar{t}$ system, $m_{t\bar{t}}$, is close to twice the top-quark mass ($m_{t\bar{t}} \approx 2 \cdot m_t$) and in the high-$m_{t\bar{t}}$ region with large scattering angles where the top quarks carry high Lorentz boosts. Therefore, unlike spin correlations, which can be measured inclusively across the $m_{t\bar{t}}$ phase space, the purity of entangled $t\bar{t}$ spin states is expected in very limited kinematic regions, making their experimental observation both novel and challenging.

These measurements are performed using different decay channels of the $t\bar{t}$ system: the dilepton channel, in which both the top quarks decay semi-leptonically into a b-jet and a W boson that subsequently decays into a charged lepton and a neutrino; the lepton+jets channel, where the W boson of one of the top quark decays hadronically into two jets, while the other top quark decays semi-leptonically; and the all-hadronic channel, in which both top quarks decay hadronically. Among these, the all-hadronic channel is the most abundant one but it is experimentally the most challenging due to the presence of six jets in the final state. Therefore, the dilepton and the lepton+jets channels are most commonly used in experimental measurements.

Motivated by these considerations, assuming SM and therefore QM, the ATLAS collaboration performed the first measurement of quantum entanglement in top-quark pairs in the threshold region [17], defined by $340 < m_{t\bar{t}} < 380\,\text{GeV}$. This measurement was carried out in the dilepton decay channel. While it represented an important milestone, with entanglement observed with a significance exceeding 5σ, discrepancies between data and theoretical predictions suggested that some aspects of the dynamics in the threshold region are not fully captured by the current modelling of the $t\bar{t}$ system. One possible explanation for these deviations is the formation of $t\bar{t}$ quasi-bound states, which arise from the onset of hadronisation of top quarks but do not fully develop due to their short lifetime. Such effects, predicted well in advance of the top-quark discovery [18–20], are expected to

occur only near the production threshold, where the top quarks have relatively low velocities. Interest in the threshold region has intensified after ATLAS [21] and CMS [22] reported an excess consistent with quasi-bound-state formation at threshold. Although such effects were not initially expected to play a significant role in measurements at the LHC, the observations highlight the remarkable level of experimental precision that has been achieved. The CMS collaboration subsequently extended studies of $t\bar{t}$ entanglement by incorporating a simplified model of such quasi-bound-state effects, both in the dilepton channel [23] and in the lepton+jets channel [24] demonstrating better agreement between data and predictions. Moreover, CMS performed the first measurement of entanglement in the lepton+jets channel in the high-$m_{t\bar{t}}$ region [24], where the top quarks are typically well separated in space, thereby extending experimental sensitivity to new kinematic regimes. This region is also expected to be sensitive to quantum states that can violate Bell inequalities, although such inequalities cannot presently be tested directly at colliders. This further reinforces the importance of understanding entanglement dynamics in this regime.

The entanglement measurement performed so far by ATLAS was carried out in the dilepton channel, where spin correlations and entanglement are probed using charged leptons, which retain a large fraction of the top-quark spin information and leave particularly clean signatures in the detector. However, this channel suffers from limited statistics. In comparison, the lepton+jets channel offers a significantly larger number of events, making it a complementary and statistically favourable channel for entanglement studies despite the existing measurement in the dilepton channel. Therefore, in this thesis, two measurements of quantum entanglement in top-quark pairs focusing on the lepton+jets channel are presented. The first focuses on the threshold region, while the second targets the high-$m_{t\bar{t}}$ regime. Both measurements use data collected by the ATLAS detector during the period 2015–2018. Unlike the first measurement, the second is an ongoing analysis effort that employs the latest software and tools available in the ATLAS framework, and is expected to eventually include data collected beyond 2018. For the threshold region, the analysis incorporates the effects of $t\bar{t}$ quasi-bound states in the predictions using a more sophisticated modelling approach, making it the first ATLAS entanglement measurement to account for such effects. Furthermore, the second measurement represents the first ATLAS entanglement measurement in the high-$m_{t\bar{t}}$ regime. Performing measurements in the lepton+jets channel presents additional challenges, particularly in identifying decay products that retain strong sensitivity to the top-quark spin. In the high-$m_{t\bar{t}}$ region, this task becomes even more challenging due to the collimation of the decay products from highly boosted top quarks. These challenges, together with the methods developed to address them, are discussed in this thesis and are essential for enabling future measurements in this statistically dominant channel.

Following an introductory overview of the SM in Chapter 2, Chapter 3 focuses on the theoretical background of quantum effects in top-quark pairs. Subsequent chapters describe the LHC and ATLAS detector (Chapter 4), the simulation of signal and background processes (Chapter 5), and the techniques used to reconstruct particle collision events in the ATLAS detector (Chapter 6). The analysis strategy and results of the measurement in the $t\bar{t}$ threshold region are presented in Chapter 7, while the corresponding results for the high-$m_{t\bar{t}}$ region are discussed in Chapter 8. The thesis concludes with a summary and outlook in Chapter 9.

Author's contributions

Experiments in particle physics probe fundamental interactions at extremely high energies, requiring large-scale facilities such as particle accelerators and colliders. To achieve these energies, complex infrastructures like the LHC are constructed, together with general-purpose detectors such as ATLAS, designed to detect and record the resulting collision data. The scale and complexity of these experiments are immense, involving advanced technologies in accelerator physics, detector design, computing, and data analysis, as well as significant financial and human resources. As a result, such projects are difficult to be realised or operated by a single institution or even a single country. Instead, they are carried out by large international collaborations composed of thousands of members. The ATLAS collaboration is one such example, consisting of physicists, engineers, technicians, students, and support staff from around the world, that make up over 5500 members. Therefore, the work presented in this thesis builds upon the contributions of many members of the ATLAS collaboration. The specific contributions of the author are summarised below.

The author is one of the two main contributors to the entanglement measurement at the $t\bar{t}$ production threshold, with primary responsibility for the development and maintenance of the statistical analysis framework. The author worked on including the $t\bar{t}$ quasi-bound-state effects within the fit model and also developing the $t\bar{t}$ reconstruction at particle level. In addition, the author studied the robustness of the statistical model and the impact of the complete set of systematic uncertainties on the measurement.

For the ongoing analysis of entanglement in the high-$m_{t\bar{t}}$ regime, the author has been involved since the early stages and has significantly shaped the analysis strategy within the team. During the initial phase, the author conducted studies on the potential use of large-R jets for reconstructing the hadronically decaying top quark, representing a possible future direction of the analysis. The author also investigated different choices of decay products that retain maximal spin information from the hadronically decaying top quark, with the aim of improving the overall sensitivity to entanglement. Furthermore, the author established the full statistical analysis framework for this measurement from the ground up, developing and validating multiple configurations that now serve as the baseline for incorporating additional components into the analysis.

In addition to the two independent analyses presented in this thesis, the author contributed to a search for $t\bar{t}$ resonances [25]. In the lepton+jets channel of the analysis, this contribution involved incorporating information from top-tagged jets, along with their associated systematic uncertainties, into the analysis framework. Furthermore, the author contributed to the production of reduced-size data samples in a format suitable for the statistical interpretation of the results.

Beyond this, the author was an active member of the ATLAS Inner Detector Tracking Group focusing on clustering and tracking in dense environments. Within this effort, the author played a leading role in the development of a framework to measure the fraction of lost tracks in dense environments, independently in data and simulation. The results of this study are a key input to the calibration of flavour-tagging algorithms used throughout the ATLAS collaboration. The measurement results are included in a recent publication on the performance of tracking algorithms [26] and will also feature prominently in a forthcoming publication on tracking performance in dense environments.

In this thesis, figures that carry the label *ATLAS* or *ATLAS preliminary* are part of ATLAS publications or public notes, respectively. Figures without an *ATLAS* label are either taken from non-ATLAS publications or from internal ATLAS documentation used for publication or internal circulation purposes. Unless explicitly stated otherwise, all figures and tables without a reference in the caption have been produced by the author.

CHAPTER 2

Fundamentals of particle physics

"Knowledge is a paradox. The more one understands, the more one realises the vastness of his ignorance."

— Arcane 2

Throughout history, scientists have sought to uncover what lies at the heart of matter and the forces that shape the universe. Many theories have been proposed, each highlighting the understanding and experimental capabilities of its time. Over the past century, a remarkable series of discoveries have revealed a rich and complex subatomic world.

The growing understanding of particles and their interactions led, in 1970s, to the establishment of a unified theoretical framework known as the ***Standard Model (SM)*** of particle physics. Today, the SM provides a consistent description of the known elementary particles, comprising the fermions, which constitute matter, and the gauge bosons, which mediate three of the four fundamental forces of nature. The theoretical foundations and structure of the SM are discussed in detail in Section 2.1. The Higgs mechanism explains electroweak symmetry breaking and the origin of particle masses within the SM framework. Section 2.1.3 describes this process in more detail and emphasises the important role of the top quark, the most massive fermion in the SM and the main focus of this work, in studies of the Higgs sector. Despite its success, the SM has certain limitations. A few of the main ones are discussed in Section 2.2.

2.1 The Standard Model of particle physics

Our current knowledge of the elementary particles that possess no substructure and are the basic building blocks of the universe, is summarised in the SM. It also describes the fundamental interactions of nature, with the exception of gravity, namely electromagnetism, the weak interaction, and the strong interaction. Although the SM does not incorporate gravitation, it has been remarkably successful in explaining phenomena at the smallest scale accessible to experiments. At these scales, the effects of gravity are negligible compared to the other fundamental interactions. In this section, the fundamental particles and interactions that form the basis of the SM are briefly introduced. The content is based on Refs. [27–29].

2.1.1 Fundamental particles

A key property that distinguishes different types of elementary particles is their intrinsic angular momentum, referred to as ***spin***. In classical mechanics, angular momentum describes the rotational motion of an object and depends on its mass, velocity, and distance from the axis of rotation. In quantum mechanics, however, angular momentum appears in two distinct forms: orbital angular momentum, associated with the motion of the particle through space, and intrinsic angular momentum, or spin, which is an inherent property of the particle itself. The existence of spin angular momentum in electrons was inferred from experiments such as the Stern-Gerlach experiment [30], in which a beam of silver atoms was passed through a non-uniform magnetic field. Instead of producing a continuous spread, the beam split into two distinct components, corresponding to two possible discrete orientations of angular momenta despite having no orbital contribution. This result provided direct evidence not only that angular momentum in quantum systems is quantised, but also for the existence of half-integer quantum numbers associated with electron spin $s = 1/2$, thereby also demonstrating the existence of non-integer quantum numbers.

The particles in the SM are therefore divided into two major categories based on their spin. Particles with half-integer spins ($s = \frac{1}{2}, \frac{3}{2}, ..$) are called ***fermions***, while those with integer spin ($s = 0, 1, 2..$) are called ***bosons***. This distinction determines how the particles interact in nature. Fermions obey the Pauli exclusion principle and follow *Fermi-Dirac statistics*, which state that no two identical fermions can occupy the same quantum state. Consequently, the stable matter that makes up the observable universe is composed of fermions. Bosons, in contrast, obey *Bose-Einstein statistics* and can occupy the same quantum state, allowing them to act as force carriers that mediate the fundamental interactions between elementary particles.

Among the bosons, ***photons*** (γ) are responsible for the electromagnetic interaction, ***gluons*** (g) mediate the strong interaction, and the $\boldsymbol{W^{\pm}}$ and ***Z bosons*** mediate the weak interaction. These particles are known as *vector bosons*, corresponding to bosons with spin 1. While photons, gluons and Z bosons are electrically neutral, the W^- (W^+) bosons carry one unit of negative (positive) electric charge. The $W^{\pm}$ and Z bosons are relatively heavy compared to other elementary particles, with masses of approximately 80.37 GeV and 91.19 GeV, respectively [31]. Photons and gluons, on the other hand, are massless. Unlike the other bosons, gluons carry colour charge and can therefore interact with themselves (subsection QCD in Section 2.1.2).

Fermions can be further divided into ***leptons*** and ***quarks***. There are six leptons, grouped into three generations with increasing mass: the ***electrons*** e, ***muons*** μ, and ***tau-leptons*** τ, along with their associated ***neutrinos*** (ν_e, ν_μ, ν_τ). Quarks come in six types, known as ***flavours***, which are grouped into three generations: ***up***(u) and ***down***(d), ***charm***(c) and ***strange***(s), ***top***(t) and ***bottom***(b). The u, c, and t quarks carry an electric charge of $+2/3$ and are collectively referred to as up-type quarks, while the d, s, and b quarks carry an electric charge of $-1/3$ and are known as down-type quarks. All fermions have corresponding ***antiparticles*** with the same mass but opposite charge-like quantum numbers. Quarks, in addition, possess a ***colour*** which leptons do not. The bosons are their own antiparticles, except for the charged W bosons.

The quarks bind together (sub-section QCD in Section 2.1.2) to form colour-neutral composite particles called ***hadrons***. These are mainly classified into ***mesons*** and ***baryons***. Mesons consist of a quark–antiquark pair that combine colour and anticolour to form a neutral state, while baryons consist of three quarks. For example, a proton is a baryon composed of two up quarks and one down quark (uud). A schematic illustration of the SM is shown in Figure 2.1.

In the following sections, the fundamental interactions are discussed in more detail (Section 2.1.2), including the mechanism responsible for the generation of particle masses, known as the Higgs mechanism (Section 2.1.3), which accounts for the existence of the spin-0 Higgs boson within the SM.

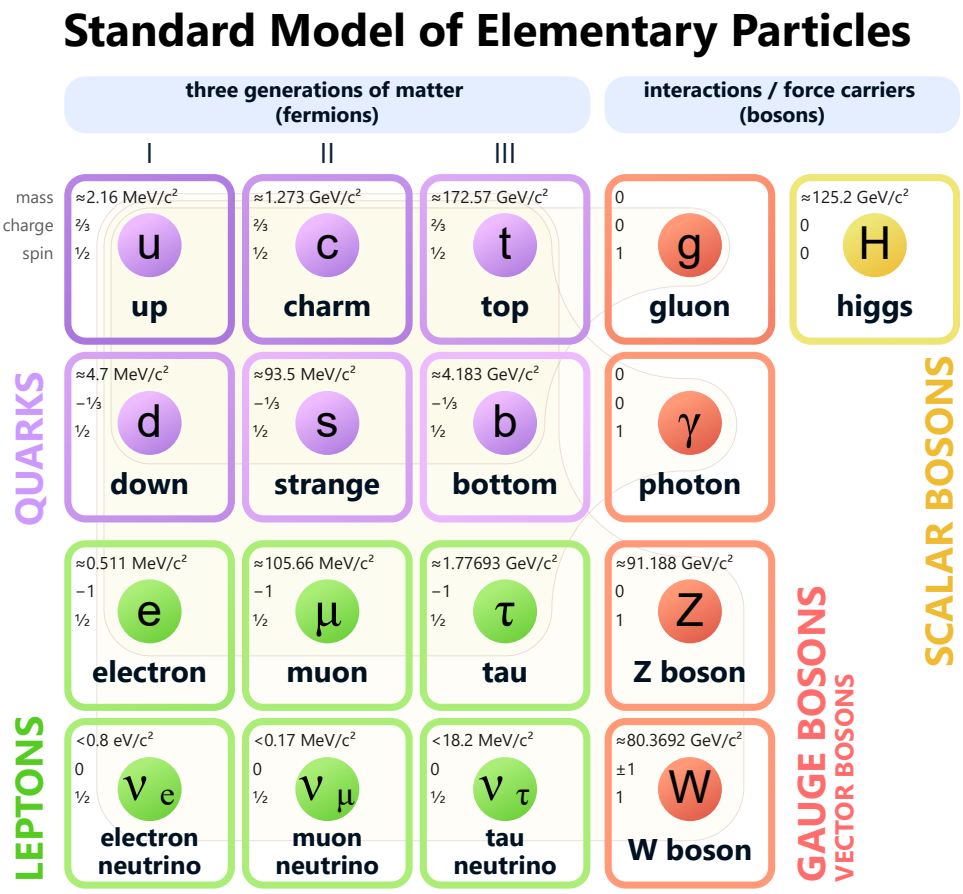


Figure 2.1: A schematic diagram of the elementary particles in the SM. The different colours indicate the fermion and boson sectors as discussed in the text. For each particle, its mass, electric charge, and spin are shown [32]. The values of the particle properties, shown in the diagram, are taken from Ref. [31].

2.1.2 Fundamental interactions

The SM is a relativistic ***quantum field theory*** that describes elementary particles and their interactions in terms of fields defined throughout space and time, denoted by $\phi_i(x)$. The dynamics of these fields are determined by the ***Lagrangian density*** $\mathcal{L}$ (often simply referred to as the Lagrangian), which is a function of the fields and their space-time derivatives, $\partial_\mu \phi_i = \frac{\partial \phi_i}{\partial x^\mu} (\mu = 0, 1, 2, 3)$. The evolution of the fields in time follows from the principle of stationary action, where the physical path of the system is obtained by extremising the action $S = \int \mathcal{L}(\phi_i(x), \partial_\mu \phi_i)\, d^4x$. This yields the Euler-Langrange

equations which provide the "equations of motion" of the fields for a given Lagrangian:

$$\partial_\mu \left(\frac{\partial \mathcal{L}}{\partial(\partial_\mu(\phi_i))} \right) - \frac{\partial \mathcal{L}}{\partial \phi_i} = 0$$

A transformation of the fields that leaves the Lagrangian unchanged is referred to as a gauge transformation. The transformed fields obey the same equations of motion, meaning that they describe the same physical system. This invariance is known as ***gauge invariance*** or *gauge symmetry*. Mathematically, these symmetries are characterised by Lie groups from group theory [33], with the relevant ones for the SM being the unitary group $U(1)$ and the special unitary groups $SU(2)$ and $SU(3)$. Any element of such a unitary group can be expressed as:

$$\hat{U} = \exp\left[i\theta_a T^a\right],$$

where T^a are the generators of the group and θ_a are real parameters. A Lagrangian that is invariant under this transformation with constant parameters θ_a is said to have *global gauge invariance*. With additional constraints, if the Lagrangian remains invariant under a transformation where the parameters are space-time dependent, $\theta_a(x)$, then it possesses *local gauge invariance*.

The SM is a ***gauge field theory*** characterised by the following product of gauge groups:

$$SU(3)_C \otimes SU(2)_L \otimes U(1)_Y, \tag{2.1}$$

which corresponds to the strong interaction between coloured (C) particles, weak interaction between left-handed (L) particles, and hypercharge (Y) interactions, respectively. The following sections discuss more closely how the fundamental interactions among elementary particles emerge naturally from the requirement of local gauge invariance and what the notations in Equation 2.1 represent.

Quantum electrodynamics (QED)

The dynamics of each of the twelve fundamental fermions, such as the electron, are described by the ***Dirac equation*** [34] of relativistic quantum mechanics. Accordingly, the Lagrangian for a free spin-$\frac{1}{2}$ particle is postulated as:

$$\mathcal{L}_{\text{Dirac}} = \bar{\psi}(i\gamma^\mu \partial_\mu - m)\psi, \tag{2.2}$$

where ψ and $\bar{\psi} = \psi^\dagger \gamma^0$ are the Dirac spinor and its adjoint, respectively, and γ^μ are the gamma matrices satisfying the Dirac anti-commutation relations $\{\gamma^\mu, \gamma^\nu\} = 2g^{\mu\nu}$, with $g^{\mu\nu}$ being the Minkowski metric tensor. While this Langrangian is invariant under a $U(1)$ global gauge transformation given by $\psi \rightarrow \psi' = [iq\theta]\psi$, it is not invariant under a $U(1)$ local gauge transformation. In the latter case, the Lagrangian acquires an additional phase-dependent term of the form $q\,\bar{\psi}\gamma^\mu(\partial_\mu\theta(x))\psi$, where q is an arbitrary constant which serves as the generator of the $U(1)$ group. To restore local gauge invariance, a new massless vector field A_μ, referred to as a gauge field, is introduced with appropriate transformation properties. This is achieved by replacing the ordinary derivative ∂_μ with the covariant derivative D_μ:

$$D_\mu = \partial_\mu + iqA_\mu,$$

where the field A_μ transforms as :

$$A_\mu \xrightarrow{U(1)} A'_\mu = A_\mu - \partial_\mu \theta(x),$$

Substituting this into the original Dirac Lagrangian of Equation 2.2 introduces an additional term, which can be interpreted as an interaction between the fermion and the gauge field.

The field theory describing the interaction of a fermion with the electromagnetic field via the exchange of gauge bosons is called ***Quantum Electrodynamics (QED)***. In this theory, the interaction term represents the coupling between photons and charged fermions with electric charge q. Therefore, the generator of the $U(1)$ symmetry group is the electric charge operator q, and the mediating gauge boson is a massless photon. For a single fermion type, the QED Lagrangian is given by:

$$\mathcal{L}_{\text{QED}} = \underbrace{\bar{\psi} i\gamma^\mu \partial_\mu \psi}_{\text{Fermion kinetic term}} - \underbrace{m\bar{\psi}\psi}_{\text{Fermion mass term}} - \underbrace{q\,\bar{\psi}\gamma^\mu \psi A_\mu}_{\text{Fermion-photon interaction term}} - \underbrace{\frac{1}{4} F_{\mu\nu}F^{\mu\nu}}_{\text{Photon field kinetic term}} ,$$

where the second term, $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$, represents the Lagrangian of the free photon field, also referred to as the kinetic term of the electromagnetic field. Including a mass term for the photon would violate gauge invariance, and therefore it is not allowed in the Lagrangian.

Quantum chromodynamics (QCD)

The quantum field theory that describes the strong interactions within the SM is called ***Quantum Chromodynamics (QCD)***. It is based on the $SU(3)_C$ symmetry group, where the subscript C denotes ***colour*** charge, the fundamental charge of QCD. Colour charge comes in three types, conventionally referred to as red, green, and blue (r, g, and b). Only quarks and gluons carry colour charge and thus participate in the strong interaction. All other bosons and leptons are colour neutral. To achieve local gauge invariance in this context, the partial derivative ∂_μ in Equation 2.2 is replaced by the covariant derivative, $D_\mu = \partial_\mu - i g_s T_a G_{a\mu}$, where T_a are the eight generators of the $SU(3)$ symmetry group, which are represented in terms of the Gell-Mann matrices as $T_a = \frac{\lambda_a}{2}$ and satisfy the commutation relation $[T_a, T_b] = i f_{abc} T_c$. Here, f_{abc} are the structure constants of $SU(3)$, g_s is the strong coupling constant, and G_a are the eight gauge fields corresponding to the eight gluons of QCD. Based on these definitions, the QCD Lagrangian for a given quark flavour can be written as:

$$\mathcal{L}_{\text{QCD}} = \underbrace{\bar{\psi} i\gamma^\mu \partial_\mu \psi}_{\text{quark kinetic term}} - \underbrace{m\bar{\psi}\psi}_{\text{quark mass term}} - \underbrace{g_s\,\bar{\psi}\gamma^\mu T_a G_{a\mu}\psi}_{\text{quark-gluon interaction term}} - \underbrace{\frac{1}{4} G_{a\mu\nu}G_a^{\mu\nu}}_{\text{Gluon field kinetic term}} ,$$

where $G_a^{\mu\nu} = \partial^\mu G_a^\nu - \partial^\nu G_a^\mu - g_s f_{abc} G_b^\mu G_c^\nu$ denotes the gluon field strength tensor. The $SU(3)_C$ group has a non-Abelian nature, meaning that the generators of the group do not commute. This implies that some of the structure constants satisfy $f_{abc} \neq 0$. As a result, the gluons themselves carry colour charge and can therefore interact with one another, leading to characteristic self-interaction terms in the gluon kinetic part of the QCD Lagrangian. This behaviour contrasts with that of the photon in QED, which carries no electric charge and thus does not couple to itself. This non-Abelian property of QCD gives rise to two key features of the theory.

Colour confinement Despite experimental efforts, quarks have never been observed as free particles[1]. This is explained by the phenomenon of ***colour confinement***, which states that coloured objects are always confined within ***colour-singlet*** states. A colour-singlet state is a combination of colour charges that is overall neutral under the $SU(3)_C$ symmetry, meaning that the total colour charge cancels out. In other words, although individual quarks carry colour charge (forming *triplet states*), such states cannot exist freely because separating a single coloured object would require infinite energy.

When two quarks are pulled apart, they continue to interact via gluon exchange. Since gluons themselves carry colour charge, these exchanges create an attractive force that compresses the colour field into a narrow flux tube between the quarks, referred to as a string. As the quarks separate, the energy stored in this flux tube increases with distance. Eventually, the energy becomes sufficient to create a new quark–antiquark pair from the vacuum, preventing isolation of the original quarks. This process continues until all quarks and antiquarks have sufficiently low energy to combine to form bound, colour-neutral states called hadrons. The most common examples are baryons, consisting of three quarks, and mesons, composed of quark-antiquark pair.

Due to colour confinement, a high-energy scattering process that produces quarks and gluons in the final state results in a cascade of hadrons, known as ***jets***, through a process called ***hadronisation***. An illustration of this process is shown in Figure 2.2. Note that this represents only one realisation of hadronisation. Other possible descriptions of the hadronisation process also exist, as discussed in Chapter 5.

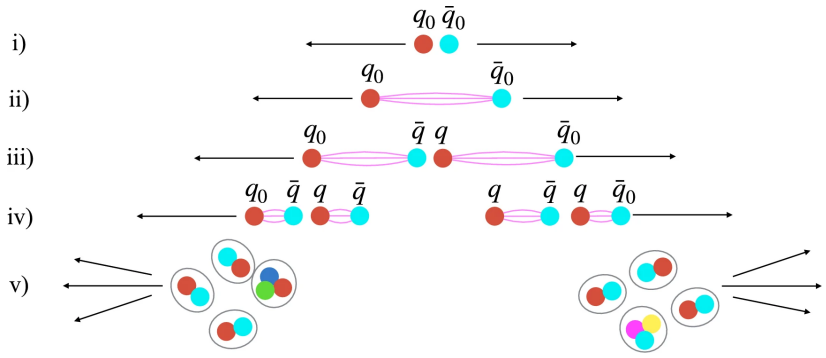


Figure 2.2: Illustration of the steps in the hadronisation process. (i) The initial quarks, q_0 and $\bar{q}_0$, begin to move apart. (ii) A narrow flux tube forms between them. (iii) The string subsequently breaks to create a new quark–antiquark pair. (iv) Repeated string breaking occurs until (v) all quarks and antiquarks have sufficiently low energy to combine into colourless hadrons. The hadronisation process results in two jets of hadrons. [35].

Running coupling and asymptotic freedom Most calculations in quantum field theory are performed using ***perturbation theory***, which expresses observables as a series expansion in powers of a coupling constant. In QCD, this coupling is the strong coupling constant, commonly expressed as $\alpha_s = g_s^2/(4\pi)$. However, some terms in this expansion become divergent due to higher-order corrections, such as Feynman diagrams[2] containing additional quark or gluon loops. A process

[1] The top quark is a special case in that it decays before it can hadronise.

[2] A Feynman diagram is a graphical depiction of interactions between subatomic particles, introduced by physicist Richard Feynman. It uses lines and vertices to show how particles move and interact, with each component representing a specific term in the mathematical equations that determine the likelihood of an interaction.

called ***renormalisation*** absorbs these divergences into an effective coupling constant, introducing a dependence of α_s on the energy scale, typically taken as the momentum transfer Q of the interaction. This dependence of the coupling on the energy scale is known as the ***running of the coupling***. A similar effect occurs in QED, where the electromagnetic coupling $\alpha = e^2/4\pi$ also depends on Q^2, though with a much weaker dependence. In QED, the coupling increases slightly with Q^2, whereas in QCD, α_s decreases with Q^2. At low energies, α_s is large, making calculations using perturbative theory difficult. In this non-perturbative regime, first-principles techniques such as lattice QCD are used for the calculations. At high energies, the decrease of α_s with Q^2 is referred to as ***asymptotic freedom***. Here, the coupling becomes sufficiently small for perturbation theory to be applicable. This behaviour naturally divides QCD into two regimes: a low-energy, non-perturbative regime, where processes like hadronisation cannot be calculated perturbatively, and a high-energy, perturbative regime, where interactions of quarks and gluons can be predicted with precision. In practical calculations, this separation is handled using factorisation, introducing factorisation and renormalisation scales (Section 3.1.1).

Electroweak theory

The construction of a gauge theory for the weak interaction that preserves local gauge invariance encounters difficulties primarily due to the large observed masses of the $W^{\pm}$ and Z bosons, which mediate this interaction. To address this issue without violating gauge symmetry, the concept of gauge invariance is extended to include both the weak and electromagnetic interactions within a unified framework known as the ***electroweak (EW) theory***. This theory is based on the symmetry group $SU(2)_L \otimes U(1)_Y$, which combines the weak isospin ($SU(2)_L$) and hypercharge ($U(1)_Y$) symmetries, as explained later. The unification of these two fundamental interactions was independently proposed by Glashow, Salam, and Weinberg, and is therefore known as the *Glashow–Salam-Weinberg (GSW) model* [36].

In the 1950s, experiments such as the Wu experiment [37] demonstrated that the weak interaction violates parity symmetry, meaning that it distinguishes between left- and right-handed fermions. It was observed that the weak interaction couples only to fermions (and antifermions) with left-(right-) handed chirality, while their right-handed counterparts do not participate in weak charged-current interactions. This is the reason for the subscript L in $SU(2)_L$. For a general Dirac field ψ, the left- and right-handed components are defined using the chiral projection operators:

$$\psi_{L/R} = \underbrace{\frac{1 \pm \gamma^5}{2}}_{\text{chiral operator}} \psi,$$

where $\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3$. These operators project out the left-(ψ_L) and right-handed (ψ_R) components of the field.

To encode this experimentally observed asymmetry, the SM assigns different transformation properties to left- and right-handed fermions. A new quantum number, called ***weak isospin*** I, and its third component I_3, are introduced to describe transformations under $SU(2)_L$. To preserve $SU(2)$ symmetry, the left-handed fermion fields are grouped into weak isospin doublets, with $I = \frac{1}{2}$ and

$I_3 = \pm\frac{1}{2}$, represented as:

$$\begin{pmatrix} \nu_e \\ e \end{pmatrix}_L, \quad \begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}_L, \quad \begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}_L, \quad \begin{pmatrix} u \\ d \end{pmatrix}_L, \quad \begin{pmatrix} c \\ s \end{pmatrix}_L, \quad \begin{pmatrix} t \\ b \end{pmatrix}_L.$$

Here, the upper and lower componenets correspond to $I_3 = +\frac{1}{2}$ and $I_3 = -\frac{1}{2}$, respectively. In contrast, the right-handed fermion fields are singlets under a $SU(2)_L$ transformation, carrying $I = 0, I_3 = 0$:

$$e_R, \mu_R, \tau_R, u_R, d_R, c_R, s_R, t_R, b_R.$$

Right-handed neutrinos are not included since they are not predicted by the SM. Since the right-handed fermions are singlets, they do not transform under $SU(2)_L$. The weak interaction thus acts only on the left-handed components, as indicated by the subscript L.

A second quantum number, referred to as the ***weak hypercharge*** Y, is defined such that:

$$Q = I_3 + \frac{Y}{2}, \tag{2.3}$$

where Q denotes the electric charge. Since the electric charge differs between the upper and lower components of an $SU(2)_L$ doublet, an $SU(2)_L$ transformation mixes states of different charge. As a result, the electromagnetic $U(1)$ symmetry of QED cannot be maintained independently in the unbroken electroweak theory. To ensure gauge invariance, a new quantum number, the weak hypercharge Y, is introduced such that both components of a given doublet share the same Y value. The $U(1)_Y$ symmetry, generated by Y, replaces the QED $U(1)$ symmetry. The familiar electromagnetic $U(1)$ symmetry is recovered after electroweak symmetry breaking, as discussed in Section 2.1.3.

To incorporate the above features, the covariant derivative of the electroweak gauge group $SU(2)_L \otimes U(1)_Y$ is defined as:

$$D_\mu = \partial_\mu + ig\frac{\sigma^a}{2}W_\mu^a + ig'\frac{Y}{2}B_\mu, \tag{2.4}$$

where the Pauli matrices denoted by $\sigma^a/2$ (with $a = 1, 2, 3$) are the three generators of the $SU(2)_L$ group. The vector triplet $\boldsymbol{W_\mu}$= $(W_\mu^1, W_\mu^2, W_\mu^3)$, associated with $SU(2)_L$, and the singlet field B_μ, associated with $U(1)_Y$, define the gauge fields of the theory. The constants g and g' are the coupling strengths corresponding to $SU(2)_L$ and $U(1)_Y$, respectively. The experimentally observed gauge bosons $W^\pm$ are obtained as linear combinations of the first two components of the $SU(2)_L$ triplet:

$$W_\mu^\pm = \frac{1}{\sqrt{2}}\left(W_\mu^1 \mp iW_\mu^2\right).$$

The remaining neutral fields, W_μ^3 and B_μ, mix to produce the physical Z boson and the photon, which mediate weak neutral and electromagnetic interactions, respectively. This mixing is described by the Weinberg rotation:

$$\begin{pmatrix} A_\mu \\ Z_\mu \end{pmatrix} = \begin{pmatrix} \cos\theta_W & \sin\theta_W \\ -\sin\theta_W & \cos\theta_W \end{pmatrix}\begin{pmatrix} B_\mu \\ W_\mu^3 \end{pmatrix},$$

where θ_W is the ***Weinberg angle*** (or weak mixing angle), defined by

$$\tan\theta_W = \frac{g'}{g}.$$

The electromagnetic coupling e is then related to the weak couplings through

$$e = g\sin\theta_W = g'\cos\theta_W.$$

The Lagrangian describing the electroweak interaction is therefore written as:

$$\mathcal{L}_{\text{EW}} = \underbrace{\bar{\psi}_L i\gamma^\mu\left(\partial_\mu + ig\frac{\sigma^a}{2}W_\mu^a + ig'\frac{Y_L}{2}B_\mu\right)\psi_L}_{\text{Left-handed fermion kinetic and interaction terms}}$$
$$+ \underbrace{\bar{\psi}_R i\gamma^\mu\left(\partial_\mu + ig'\frac{Y_R}{2}B_\mu\right)\psi_R - \frac{1}{4}W^a_{\mu\nu}W^{a\,\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu}}_{\text{Right-handed and gauge field kinetic/self-interaction terms}}.$$

Here, the non-Abelian and Abelian field-strength tensors are defined as:

$$W^a_{\mu\nu} = \partial_\mu W^a_\nu - \partial_\nu W^a_\mu + g\,f^{abc}W^b_\mu W^c_\nu,$$
$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu.$$

It is evident from the form of the Lagrangian in Equation 2.1.2 that the gauge bosons corresponding to the $SU(2)_L$ and $U(1)_Y$ fields are massless, as required by local gauge invariance. Experimentally, however, the photon is massless while the $W^\pm$ and Z bosons possess large masses. Introducing mass terms for these gauge bosons by hand would explicitly break gauge invariance and render the theory non-renormalisable. This inconsistency is resolved through the Higgs mechanism, as a consequence of which the $W^\pm$ and Z bosons acquire mass while preserving local gauge invariance.

2.1.3 The Higgs mechanism

To introduce mass terms for the $W^\pm$ and Z bosons while preserving gauge invariance and renormalisability, a new mechanism, proposed independently in 1964 by Brout and Englert [38], Higgs [39], and Guralnik, Hagen, and Kibble [40], was developed. This mechanism, now known as the ***Brout–Englert–Higgs mechanism*** or simply the *Higgs mechanism*, achieves mass generation through ***spontaneous symmetry breaking (SSB)*** of the electroweak symmetry. This means that although the full Lagrangian remains invariant under the $SU(2)_L \otimes U(1)_Y$ gauge group, the ground (vacuum) state does not share this symmetry.

To implement this mechanism, a weak isospin doublet of complex scalar fields with hypercharge $Y = +1$ is introduced:

$$\Phi = \begin{pmatrix}\phi^+\\ \phi^0\end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix}\phi_1 + i\phi_2\\ \phi_3 + i\phi_4\end{pmatrix}.$$

The Lagrangian for the scalar field is given by

$$\mathcal{L}_{\text{Higgs}} = (D^{\mu}\Phi)^{\dagger}(D_{\mu}\Phi) - V(\Phi),$$

where D_{μ} is the electroweak covariant derivative defined in Equation 2.4, and $V(\Phi)$ is the scalar potential,

$$V(\Phi) = \mu^2\,\Phi^{\dagger}\Phi + \lambda\,(\Phi^{\dagger}\Phi)^2, \qquad \text{with } \lambda > 0 \text{ and } \mu^2 < 0.$$

For $\mu^2 < 0$, the potential takes the shape of a "Mexican hat", as illustrated in Figure 2.3. The potential has infinitely many degenerate minima. Without loss of generality, the vacuum can be chosen along the neutral component:

$$\langle\Phi\rangle_0 = \frac{1}{\sqrt{2}}\begin{pmatrix}0\\ v\end{pmatrix}, \qquad v = \sqrt{-\frac{\mu^2}{\lambda}}.$$

Here, v is the ***vacuum expectation value (VEV)*** of the potential. The choice of vacuum breaks the symmetry of the Lagrangian, a process known as spontaneous symmetry breaking.

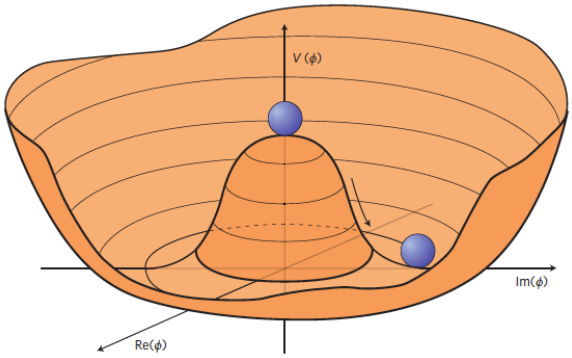


Figure 2.3: An illustration of the Higgs potential in lower dimensions, where ϕ is taken to be a single complex scalar field rather than an $SU(2)$ doublet. The parameters satisfy $\mu^2 < 0$ and $\lambda > 0$, leading to the characteristic "Mexican hat" shape of the potential [41].

Fluctuations around this vacuum correspond to the physical excitations of the field. These can be obtained by considering perturbations of the field around the vacuum state. Four degrees of freedom for perturbations relative to the chosen vacuum state exist. After spontaneous symmetry breaking, three of these become massless scalar excitations known as *Nambu–Goldstone bosons*. A gauge transformation, referred to as the unitary gauge, removes these massless scalar fields from the Langrangian. The degrees of freedom associated to these are absorbed by the $W^{\pm}$ and Z gauge fields, providing them with longitudinal components and giving them mass. The remaining degree of freedom manifests as a real scalar field, the physical ***Higgs boson***, described by the field $h(x)$.

In the unitary gauge, the field can be written as:

$$\Phi(x) = \frac{1}{\sqrt{2}}\begin{pmatrix}0\\ v + h(x)\end{pmatrix},$$

where $h(x)$ represents fluctuations about the vacuum expectation value. The symmetry breaking

pattern is therefore

$$SU(2)_L \otimes U(1)_Y \longrightarrow U(1),$$

where $U(1)$ is the gauge group of QED. The photon, therefore, remains massless while the W and Z bosons acquire mass. The Higgs sector of the Langrangian can be written as:

$$\begin{aligned}\mathcal{L}_{\text{Higgs}} = & \underbrace{\tfrac{1}{2}\partial_\mu h\,\partial^\mu h - \lambda v^2 h^2}_{\text{Higgs kinetic and mass terms}} + \underbrace{\frac{1}{2}\left(\frac{g^2 v^2}{4} W_\mu^+ W^{-\mu} + \frac{v^2(g^2+g'^2)}{4} Z_\mu Z^\mu\right)}_{\text{Mass terms of the } W^\pm \text{ and } Z \text{ bosons}} \\ & + \underbrace{\frac{g^2 v}{2} h\, W_\mu^+ W^{-\mu} + \frac{g^2}{4} h^2\, W_\mu^+ W^{-\mu} + \frac{(g^2+g'^2)v}{4} h\, Z_\mu Z^\mu + \frac{(g^2+g'^2)}{8} h^2\, Z_\mu Z^\mu}_{\text{Couplings of the Higgs field to } W^\pm \text{ and } Z \text{ bosons}} \\ & - \underbrace{\lambda v h^3 - \tfrac{\lambda}{4} h^4}_{\text{Higgs trilinear and quartic self-couplings}} + \text{constant}.\end{aligned} \tag{2.5}$$

From the above equation, masses of the electroweak gauge bosons are obtained as:

$$m_{W^\pm} = \frac{1}{2} v g,\; m_Z = \frac{1}{2} v \sqrt{g^2 + g'^2},\; \text{and}\; m_A = 0,$$

where the photon remains massless as expected from the unbroken U(1) gauge symmetry. The fluctuation of the Higgs field around the vacuum state corresponds to the physical scalar Higgs boson, with mass $m_h = \sqrt{2\lambda} v$. The terms cubic and quartic in the Higgs field represent self interactions between Higgs bosons. Equation 2.5 additionally describes interaction vertices of the form WWH, ZZH, $WWHH$, and $ZZHH$. These correspond to couplings of a single Higgs boson to two electroweak vector bosons, as well as quartic couplings involving two Higgs bosons and two electroweak vector bosons.

The discovery of the Higgs boson was announced on the 4$^{\text{th}}$ of July 2012 at the CERN LHC where the ATLAS [10] and CMS [11] experiments (Chapter 4) observed a "Higgs-like" particle at a mass of approximately 125 GeV [41].

Fermion masses The Higgs mechanism can be extended to generate fermion masses through interactions between the fermion fields and the scalar Higgs field, known as Yukawa interactions. These take the form:

$$\mathcal{L}_Y = -y_e \left(\bar{\psi}_L \Phi \psi_R + \bar{\psi}_R \Phi^\dagger \psi_L\right) = -\frac{y_e}{\sqrt{2}}(v+h)\left(\bar{\psi}_L \psi_R + \bar{\psi}_R \psi_L\right),$$

where y_e is the ***Yukawa coupling*** constant, a free parameter of the model.

The fermion mass term arises from the vacuum expectation value of the Higgs field:

$$m_f = \frac{y_e v}{\sqrt{2}}.$$

After spontaneous symmetry breaking, a Higgs–fermion interaction term of the form

$$\mathcal{L}_{hff} = -\frac{m_f}{v} h\bar{\psi}\psi$$

remains, indicating that the Higgs–fermion coupling is proportional to the fermion mass. As a result, heavier fermions interact more strongly with the Higgs field. In particular, the top quark, being the most massive known fermion, exhibits the largest Yukawa coupling and plays a crucial role in studies of effects in the Higgs sector that go ***beyond the SM (BSM)***.

After introducing fermion masses, the same mechanism can be extended to the quark sector. Unlike leptons, quarks appear in three generations, and each type (up-type and down-type) possesses its own set of Yukawa couplings. When the Higgs field acquires its vacuum expectation value, these couplings give rise to mass matrices for the up- and down-type quarks. However, these mass matrices are, in general, not diagonal, meaning that the quark fields that participate in weak interactions, known as *weak eigenstates*, are not identical to the *mass eigenstates* that arise after SSB. The fields to which the $W^{\pm}$ bosons couple form isospin doublets that are eigenstates of the weak isospin operator rather than the mass operator. As a result, interactions mediated by the charged electroweak bosons are flavour-changing, allowing transitions between different quark generations. In the quark sector, the relation between the weak and mass eigenstates is described by the ***Cabibbo–Kobayashi–Maskawa (CKM) matrix*** [42, 43], a 3×3 unitary matrix parameterised by three mixing angles (θ_{12}, θ_{23}, θ_{13}) and a CP[3]–violating phase (δ_{13}). The magnitudes of the CKM matrix elements are determined experimentally and can be expressed as:

$$\begin{pmatrix} |V_{ud}| & |V_{us}| & |V_{ub}| \\ |V_{cd}| & |V_{cs}| & |V_{cb}| \\ .|V_{td}| & |V_{ts}| & |V_{tb}| \end{pmatrix} \approx \begin{pmatrix} 0.974 & 0.225 & 0.004 \\ 0.225 & 0.973 & 0.041 \\ 0.009 & 0.040 & 0.999 \end{pmatrix}. \tag{2.6}$$

Each element V_{ij} represents the strength of the weak transition between an up-type quark of flavour i and a down-type quark of flavour j. The magnitudes of these elements determine the relative probabilities of flavour-changing processes in the quark sector and thus play a crucial role in quark decays mediated by the charged weak current.

For the bottom (b) quark, weak decays occur predominantly via transition to up-type quarks , such as $b \to cW^-$ or $b \to uW^-$ (it cannot decay to a top quark since the top quark is heavier). The probabilities of these processes are proportional to $|V_{cb}|^2$ and $|V_{ub}|^2$, respectively. Since $|V_{cb}| >> |V_{ub}|$, decays involving charm quarks dominate. However, since $|V_{cb}|^2$ is still significantly smaller than unity, the weak decay of the b quark proceeds more slowly than many other quark decays, resulting in bottom hadrons having relatively long lifetimes. These distinctive decay signatures are exploited experimentally to identify jets originating from b-quarks using dedicated algorithms (Section 6.4.6)

In the case of the top (t) quark , the situation is simpler. It decays almost exclusively via $t \to W^+b$, accounting for more than 99% of top-quark decays [31]. This decay process is governed by the CKM element V_{tb}. Since this element is very close to unity, decays to other quark flavours are heavily suppressed.

[3] Here, C stands for charge conjugation and P for parity.

2.2 Limitations of the Standard Model

Although the SM provides a successful description of the fundamental particles and their interactions, it leaves certain questions unanswered. This section highlights some of the key limitations of the SM.

Inclusion of gravity An accurate classical description of gravitation is provided by the theory of general relativity. However, since this theory is not a quantum theory and applies primarily to phenomena at large scales, no comprehensive framework exists that reconciles it with the quantum field theory underlying the SM [44]. The SM is valid only up to energies approaching the *Planck scale,* $O(10^{19}\,\text{GeV})$, beyond which the effects of gravity become too significant to be neglected.

Matter-antimatter asymmetry In current cosmological models, equal amounts of matter and antimatter were produced during the early stages of the universe. However, the present universe is dominated by matter. This observed imbalance is referred to as the *matter-antimatter asymmetry*. To explain this phenomenon, Sakharov proposed a set of necessary conditions, known as *Sakharov conditions* [45], that must be satified for this asymmetry to develop. These conditions require baryon number violation, C- and CP-symmetry violation, and interactions out of thermal equilibrium. While all three conditions are present to some extent within the SM [46, 47], their effects are insufficient, within understanding, to account for the magnitude of the observed matter-antimatter imbalance in the universe.

Dark matter and dark energy Astrophysical observations, beginning with measurements of *rotational velocities of spiral galaxis* in the 1970s by Vera Rubin [48], followed by the detection of anisotropies in the *Cosmic Microwave Background (CMB)* [49], and *gravitational lensing* [50], indicate that the majority of the matter content of the universe consists of a form of matter that does not participate in the electromagnetic interactions. This non-luminous component is known as *dark matter*. Furthermore, the observed accelerated expansion of the universe is attributed to a form of energy referred to as *dark energy*. Together, dark matter and dark energy are estimated to constitute about 95% of the mass-energy content of the universe [51]. While various BSM theories propose potential dark matter candidates, such as *heavy sterile neutrinos* [52], *axions* [53], and *weakly interacting massive particles (WIMPS)* [54], the nature and origin of dark energy remains unclear within the framework of the SM.

Neutrino masses Analogous to the CKM matrix (Equation 2.6) in the quark sector, the *Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix* [55] describes flavour mixing among neutrinos. It demonstrates that neutrino flavour and mass eigenstates are not identical and implies that neutrinos must possess non-zero masses, as experimentally shown by the observation of *neutrino oscillations* [56, 57]. However, within the original formulation of the SM, the neutrinos are predicted to be massless. Extending the SM to incorporate finite neutrino masses naturally raises the question of whether neutrinos are Dirac or *Majorana* particles [58]. Determining which of these descriptions, if either, is correct, remains one of the open problems.

The electroweak hierarchy problem Another problem of the SM related to particle masses concerns the mass of the Higgs boson itself. The Higgs mass receives contributions from quantum corrections

involving loops of massive particles. These corrections diverge quadratically with the momentum scale. Unlike the masses of fermions and vector bosons, the physical Higgs mass is not protected against such divergent contributions. At the Planck scale ($\mathcal{O}(10^{19}\,\text{GeV})$), up to which the SM is assumed to be valid and at which quantum gravitational effects are expected to become significant, the loop corrections to the Higgs mass would naturally be comparable to this scale. To obtain the observed Higgs mass of $m_H \approx 125\,\text{GeV}$, an extremely precise fine-tuning of the bare Higgs mass would be required. This severe sensitivity to high-energy physics makes the theory appear unnatural, giving rise to the so-called *hierarchy problem*. The ratio between the *electroweak scale (*$\mathcal{O}(100\,\text{GeV})$*)* and the Planck scale is approximately 10^{17}, representing an enormous separation between "low energy" physics (SM) and the "highest energy" scale (Planck). This huge gap is precisely why it is referred to as the hierarchy problem. Many BSM theories, such as models with an extended Higgs sector [59] or supersymmetry [60], have been proposed to address this issue, typically introducing new particles or interactions at the TeV scale that stabilise the Higgs mass against large quantum corrections.

CHAPTER 3

Top quarks: probing the intersection of particle physics and quantum theory

"If you think you understand quantum mechanics, you don't understand quantum mechanics."

— Richard Feynman

As the heaviest known elementary particle in the SM, with a mass of 172.69 ± 0.30 GeV [31], the top quark holds a unique significant position among all fundamental particles. Its strong coupling to the Higgs boson links it directly to the mechanism of electroweak symmetry breaking. Consequently, a deeper understanding of its properties can provide valuable insights into fundamental particle interactions and shed light on the broader workings of the universe.

Unlike lighter quarks, the top quark has an extremely short lifetime ($\sim 10^{-25}$s) [61], which is shorter than the typical hadronisation ($\sim 10^{-23}$s) [62] and spin decorrelation ($\sim 10^{-21}$s) [31] timescales. As a result, it decays before it can hadronise, preventing the formation of top-flavoured hadrons. This unique property allows physicists to study the top quark as a "bare" quark, offering a rare opportunity to probe the behaviour and, in particular, the spin properties of an unconfined quark directly.

One of the most important consequences of this, is the ability to access the spin of the top quark through its decay products. Spin, as introduced in Section 2.1, is an intrinsic form of angular momentum carried by elementary particles, and it plays a fundamental role in determining how particles interact. At the LHC, where top quarks are predominanly produced in pairs, a unique opportunity emerges to probe effects such as spin correlations and entanglement in top-quark pairs. In contrast to low-energy tests of QM using photons, ions, or electrons [12–16], high-energy proton–proton (pp) collisions enable direct investigation of these quantum effects in the quark sector and at unprecedented energy scales. This paves the way for studying quantum aspects at high-energy colliders, thus offering a novel bridge between quantum theory and particle physics.

This chapter begins by introducing the production and decay mechanisms of top-quark pairs at the LHC (Section 3.1). Section 3.2 then delves into the conceptual foundations and experimental strategies used to study quantum effects, specifically spin correlations, quantum entanglement, and Bell correlated states, and how these can be accessed through the kinematic properties of the top-quark

decay products. The chapter also discusses the key challenges and limitations involved in measuring such effects in a hadron collider environment, and highlights the various approaches developed to address them. It concludes with a discussion of potential $t\bar{t}$ quasi-bound-state effects (Section 3.3), which have recently attracted interest and have recently been observed in the low $m_{t\bar{t}}$ region, which is one of the regions where entanglement is expected to be measured.

3.1 Production and decay mechanisms of top quarks in the SM

In the SM, top quarks at the LHC are primarily produced through both strong and electroweak interactions, with the dominant contribution arising from strong-interaction processes, most notably ***top-antitop pair ($t\bar{t}$) production***, which forms the focus of the entanglement measurements presented in this thesis. In addition, top quarks can also be produced singly via electroweak interactions through the t-, s-, and associated tW associated production modes, collectively referred to as *single-top production* (Section 5.4.2). Due to its short lifetime, the top quark decays almost exclusively via the weak interactions, predominantly into a W boson and a bottom quark. A more detailed discussion of the top-quark production and decay mechanisms can be found in Ref. [31].

3.1.1 Top-quark pair production

The rate of production processes, such as $t\bar{t}$ production, is quantified by the ***cross-section***, a physical observable that depends on the underlying particle interactions and is proportional to the probability for a given process to occur as a result of pp collisions. Cross-sections can be predicted theoretically and measured experimentally, and are expressed in units of area, conventionally in *barns* (1 barn = $10^{-28}\mathrm{m}^2$). In quantum field theory, a cross-section is calculated by first identifying all Feynman diagrams that contribute to the process at a given order in perturbation theory. Each diagram represents a possible interaction channel and is constructed from the relevant Lagrangian densities discussed in Section 2.1.2. These diagrams are translated into a mathematical expression called the *scattering amplitude* or the *matrix element* $\mathcal{M}$. All the contributing diagrams are added and the measurable quantity is given by the squared matrix element $|\mathcal{M}|^2$, summed over final-state quantum numbers and averaged over initial ones. $|\mathcal{M}|^2$ is then integrated over the appropriate phase space and normalised by the incident flux.

At the LHC, when protons collide, it is actually their constituents- quarks and gluons, collectively known as ***partons***, that interact with each other. This primary interaction which occurs over extremely short distances and at high energies is referred to as the ***hard scattering process***. In the SM, $t\bar{t}$ production in pp collisions is governed by hard scatterings between partons i and j of the colliding protons, occurring at a momentum transfer scale Q where *perturbative Quantum Chromodynamics (pQCD)* is applicable.

However, since quarks and gluons are confined within protons, their momenta are not fixed. This long-distance, non-perturbative aspect of the structure of the proton is expressed in terms of ***Parton Distribution Functions (PDFs)***, $f_i(x, Q^2)$, which give the probability of finding a parton of type i carrying a fraction x of the momentum of the proton at the scale Q.

The separation between these two regimes: short-distance perturbative physics, which corresponds to the hard-scattering process, and long-distance non-perturbative physics, encoded in the PDFs, is formalised through the ***factorisation theorem*** [63]. In order to factorise the hard-scattering

component from the long-distance one, two energy scales are defined. The first one is known as the ***factorisation scale*** μ_F, which separates perturbative from non-perturbative physics. The second one is the ***renormalisation scale*** μ_R, which arises from a choice of a renormalisation scheme to remove low-energy divergences from the finite physical objects. According to this principle, the total $t\bar{t}$ (plus anything else, denoted by X) production cross-section in pp collisions at a centre-of-mass energy $\sqrt{s}$ (see Section 4.1.1 for its definition) can be factorised into a convolution of the PDFs with the hard-scattering component, as follows [63] :

$$\sigma_{pp\to t\bar{t}+X}(s, m_t) = \sum_{i,j} \int_{4m_t^2}^{s} d\hat{s} \mathcal{L}_{ij}(s, \hat{s}, \mu_F)\, \hat{\sigma}_{ij}(\hat{s}, m_t, \mu_R, \mu_F). \tag{3.1}$$

Here, $\hat{s}$ denotes the partonic centre-of-mass energy squared. The term $\hat{\sigma}_{ij}$ represents the partonic cross-section, which depends on $\hat{s}$, the top quark mass m_t, the strong coupling constant α_s, and the renormalisation and factorisation scales. The lower limit of the integral reflects the kinematic requirement that the partonic system must provide at least $2m_t$ of energy to produce the two top quarks at rest in their centre-of-mass frame, while the upper limit corresponds to the maximum partonic energy, $\hat{s}$, attainable in a pp collision.

While the physical cross-section is, in principle, independent of the choice of the scales μ_R and μ_F, perturbative calculations are performed only up to a finite order. This introduces a residual dependence on the scales, which can be minimised by chosing μ_R and μ_F to be around the characteristic energy scale of the process. Variations around this choice are typically used to estimate the theoretical uncertainty on the cross-section.

The parton luminosity function $\mathcal{L}_{ij}$ provides the probability of finding partons i and j inside the protons with sufficient energy to produce a system with centre-of-mass energy $\sqrt{\hat{s}}$. It is computed from the PDFs of the incoming partons. These PDFs are extracted from a global fit on a wide range of experimental data. The high-Q^2, low-x regime of deep inelastic scattering was extensively studied at the HERA electron-proton collider at DESY [64]. Since then, measurements from $p\bar{p}$ and pp collisons at the Tevatron and the LHC, respectively, have extended the kinematic coverage to higher values of x. To achieve the best agreement with experimental data, the PDFs are evolved according to constraints imposed by the theoretical framework of QCD, such as the DGLAP (Dokshitzer-Gribov-Lipatov-Altarelli-Parisi) equations [27]. The output of this procedure is a set of PDFs at a given Q^2 scale. As an example, the PDF4LHC PDF sets [65] are displayed in Figure 3.1 as a function of the incoming proton momentum fraction x. Since the proton contains valence u- and d-quarks (denoted by u_V and d_V), the probability of finding them is significantly larger than those of the corresponding antiquarks, particularly at higher momentum fractions ($x > 0.1$). In this region, the momentum of the proton is carried primarily by the ***valence quarks***, while antiquarks remain suppressed. In contrast, antiquarks originate predominantly from gluon splitting processes ($g \to q\bar{q}$) and form the so-called ***sea*** quarks, which are concentrated at low values of x. At small x ($x < 0.1$), the gluon distribution dominates over the individual quark contributions. This dominance arises because gluons multiply rapidly through successive splitings ($g \to gg$), leading to a steep rise in their density compared to that of quarks.

The $t\bar{t}$ production cross-section can be calculated using pQCD as an expansion in α_s. At ***leading order (LO)*** in this pQCD expansion, the dominant contributions come from gluon-gluon (gg) fusion and quark-antiquark ($q\bar{q}$) annihilation processes. At this order, the cross-sections are proportional to α_s^2. Figure 3.2 shows the corresponding Feynman diagrams for SM $t\bar{t}$ production, which are $2 \to 2$

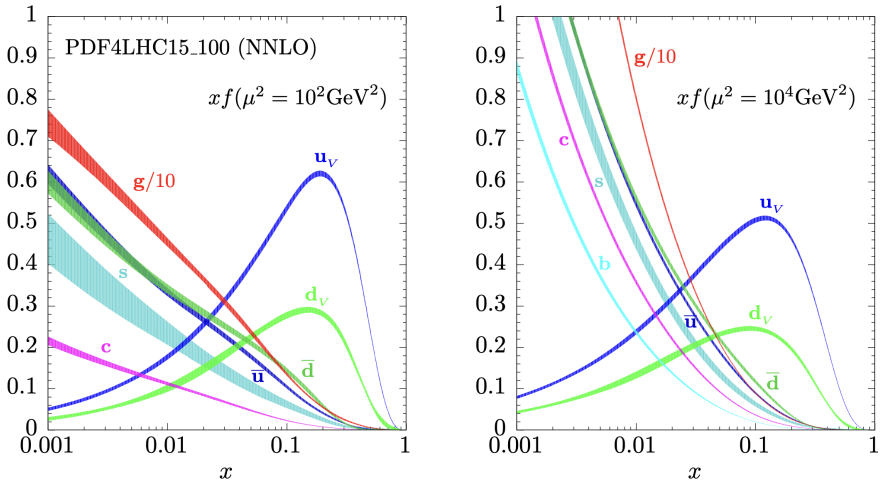


Figure 3.1: PDF4LHC PDF sets [65] as a function of x, the proton momentum fraction, for different choices of the factorisation scale, denoted by μ in the figures [66]. The left plot corresponds to a factorisation scale of 10^2 GeV^2, while the right one corresponds 10^4 GeV^2. The y-axis shows the quantity $xf(x, \mu^2)$, such that the distributions represent the fraction of the proton momentum carried by different partons.

processes with two partons in the initial and two in the final state. Due to the steep rise of the gluon PDF at low momentum fraction x, as shown in Figure 3.1, gg fusion overwhelmingly dominates at LHC energies.

At ***next-to-leading order (NLO)***, the cross-section receives corrections proportional to α_s^3. These include real-emission $2 \to 3$ processes, where an additional parton is emitted, such as $gg \to t\bar{t}g$, as well as contributions from initial states involving gq and $g\bar{q}$. Virtual loop diagrams also modify the LO $2 \to 2$ processes. Such corrections account for additional QCD effects and improve the accuracy of the theoretical prediction.

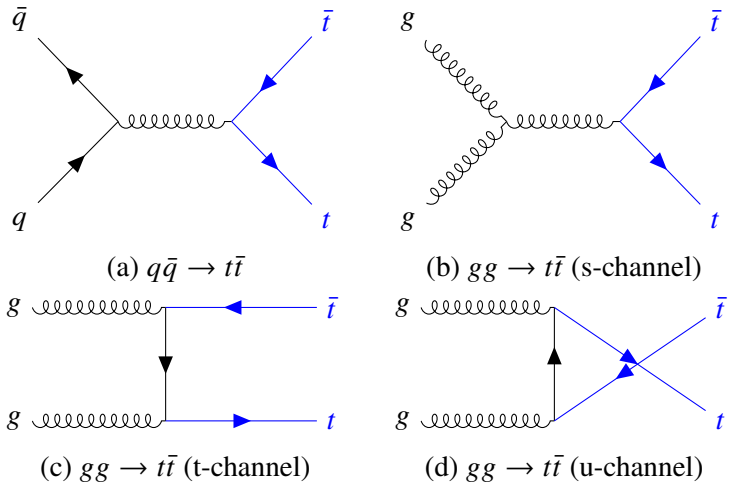


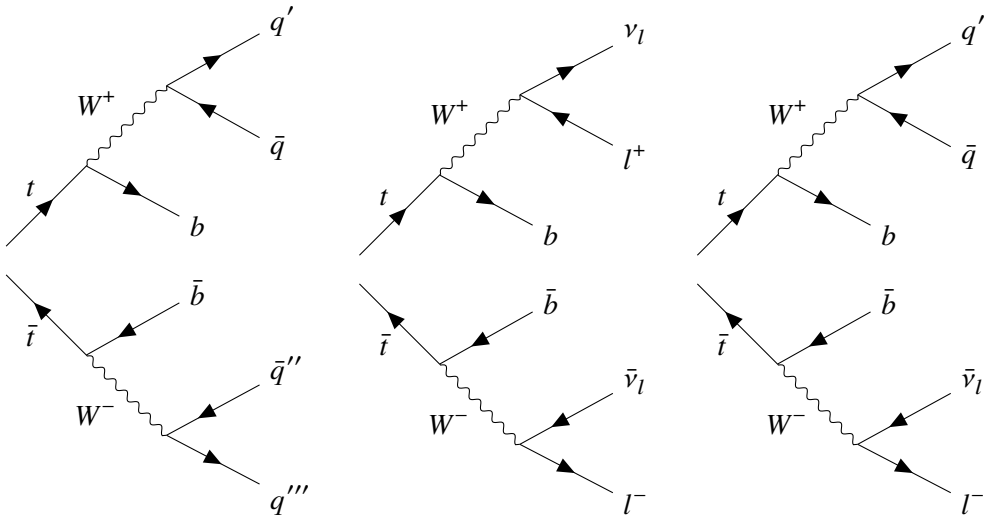
Figure 3.2: LO Feynman diagrams for $t\bar{t}$ production: (a) quark–antiquark annihilation and (b–d) gluon fusion via s-, t-, and u-channels.

3.1.2 Top-quark decays

Top quarks decay exclusively via weak interactions into a W boson and down-type quark (d,s,b). Due to the structure of the CKM matrix (Section 2.1.3), the decay to the lighter down-type quarks are highly suppressed, resulting in the decay $t \to Wb$. The W boson can further decay either leptonically, into a charged lepton and a neutrino, or hadronically, into a pair of quarks that hadronise to form jets. Based on this, the decay of a $t\bar{t}$ pair can be classified into three categories, also referred to as decay channels or modes:

- **All-hadronic channel**: Both W bosons decay hadronically, forming jets.
- **Dileptonic channel**: Both W bosons decay leptonically.
- **Lepton+jets channel**: One W boson decays leptonically and the other hadronically.

Feynman diagrams illustrating these decay channels are shown in Figure 3.3.



(a) All-hadronic (b) Dileptonic (c) Lepton+jets

Figure 3.3: Feynman diagrams illustrating the three different $t\bar{t}$ decay channels.

Among these, the all-hadronic channel has the highest branching fraction (BF) of 45%. However, due to the presence of six jets in the final state, this channel is the most challenging to work with experimentally. Substantial combinatorial issues arise from limited ability to identify the flavour and origin of jets, leading to incorrect assignments of jets to the two top quarks. The dilepton channel, on the other hand, yields fewer events compared to the all-hadronic channel with a BF of 10.5%. However, the requirement of two leptons in the final state significantly suppresses the combinatorial background. Despite this advantage, reconstructing the full event remains challenging due to the presence of two neutrinos, which cannot be directly detected. As a result, various sophisticated reconstruction techniques such as neutrino weighting [67] are employed to determine their respective four-momenta. Most analytical methods impose constraints on the W-boson and top-quark masses, allowing a system of equations to be solved and the full neutrino momenta to be determined.

This work focuses on measurements in the lepton+jets channel. This channel has a higher BF of 43.8% compared to the dilepton channel, making it statistically more advantageous. With only one neutrino in the final state, the event kinematics are better constrained than in the dilepton channel. Moreover, it offers a lower combinatorial background compared to the all-hadronic channel. However, the measurement presented in this work still heavily relies on the accurate assignment of jets in the final state, which remains one of the main challenges of this channel. The BFs of the three decay modes are listed in Table 3.1.

Decay channel	Final state	Branching fraction (%)
All-hadronic	$(q'\bar{q}b)(q'''\bar{q}''\bar{b})$	45.7
Dileptonic	$(l^+\nu_l b)(l^-\bar{\nu}_l\bar{b})$	10.5
Lepton+jets	$(l^+\nu_l b)(q'''\bar{q}''\bar{b})$ or $(q'\bar{q}b)(l^-\bar{\nu}_l\bar{b})$	43.8

Table 3.1: Branching fractions of the different $t\bar{t}$ decay channels, as reported in [31]. Here, all flavours of leptons are included $l = e, \mu, \tau$.

3.2 Quantum properties of top quarks

Top quarks, due to their extremely short lifetimes, fulfill the requirements to experimentally probe fundamental properties of QM in a quark-antiquark system. It is the intrinsic spin of the top quark that plays a central role in these studies, acting as the carrier of quantum information. Being a fermion, the top quark possesses two possible spin states, $+\frac{1}{2}$ and $-\frac{1}{2}$, commonly referred to as "spin up" and "spin down", respectively. This binary nature makes it a two-level quantum system, or a *qubit* in the language of quantum information theory. As such, its spin state can be described using the same mathematical formalism applied to qubits and can be inferred from its decay products due to its short lifetime. This correspondence allows tools from quantum information theory to be employed in studying quantum effects in $t\bar{t}$ final states.

In the framework of QM, a quantum state is described by a ***density matrix*** $\boldsymbol{\rho}$, a non-negative operator in Hilbert space, the mathematical space of quantum states, satisfying tr(ρ)= 1. For a bipartite (two-part) Hilbert space formed by the direct product of two subsystems, a and b, a quantum state is said to be separable if it can be written as

$$\rho = \sum_{n=1}^{N} p_n \rho_n^a \otimes \rho_n^b, \sum_{n=1}^{N} p_n = 1, p_n \geq 0. \tag{3.2}$$

where ρ_n^a and ρ_n^b are quantum states of the subsystems a and b, and n labels the different product states in the mixture, each weighted by a probability p_n. Any classically correlated state in the Hilbert space can be expressed in this form [68]. The special case $N = 1$ and $p_n = 1$ corresponds to a pure product state. In the case of the $t\bar{t}$ system, one deals with a ***mixed state***.

Using this formalism, the simplest model of a bipartite quantum system is a pair of qubits, with each qubit corresponding to the spin-$\frac{1}{2}$ nature of e.g a top or antitop quark. The combined system has a Hilbert space of dimension 2×2. The spin density matrix ρ of this state can be written in the

standard basis of Pauli matrices as:

$$\rho = \frac{\mathbb{I}_4 + \sum_{i=1}^{3} \left(B_i^+ \sigma^i \otimes \mathbb{I}_2 + B_i^- \mathbb{I}_2 \otimes \sigma^i \right) + \sum_{i,j=1}^{3} C_{ij} \sigma^i \otimes \sigma^j}{4}, \tag{3.3}$$

where $\mathbb{I}_n$ is the $n \times n$ identity matrix, σ^i are the corresponding Pauli matrices, $B_i^+ = \langle \sigma^i \otimes \mathbb{I}_2 \rangle$ and $B_i^- = \langle \mathbb{I}_2 \otimes \sigma^i \rangle$ denote the expectation values of the individual ***spin polarisations***[1] of the top and antitop quarks along the i-th axis in the chosen spin basis, and $C_{ij} = \langle \sigma^i \otimes \sigma^j \rangle$ encodes the ***spin correlations*** between them. As a result, by measuring all these 15 expectation values, $\langle \sigma^i \otimes \mathbb{I}_2 \rangle$, $\langle \mathbb{I}_2 \otimes \sigma^i \rangle$, $\langle \sigma^i \otimes \sigma^j \rangle$, the quantum state can be experimentally reconstructed. Measuring all components individually provides a stringent test of the SM.

By incorporating an analogous spin density matrix into the definition of the $t\bar{t}$ production cross-section, as given in Equation (3.1), and considering the angular differential cross-section of the decay products, the $t\bar{t}$ production process can be expressed as:

$$\frac{1}{\sigma} \frac{d\sigma}{d\Omega_1 \Omega_2} = \frac{1}{4\pi^2} \left(1 + \alpha_1 \mathbf{B}_1 \cdot \hat{e}_1 + \alpha_2 \mathbf{B}_2 \cdot \hat{e}_2 + \alpha_1 \alpha_2 \hat{e}_1 \cdot \mathbf{C} \cdot \hat{e}_2 \right), \tag{3.4}$$

where $\hat{e}_1$ and $\hat{e}_2$ are the directions of top (antitop) quark decay products in each of the top (antitop) rest frames, and $\Omega_{1,2}$ represent the corresponding solid angles of the top (antitop) quark decay products. The vectors $\mathbf{B}^{\pm}$ and the matrix $\mathbf{C}$ describe the top/antitop spin polarisations and the spin correlation matrix, respectively. The quantities α_1 and α_2 are called the spin analysing powers of the respective top quark decay products, defined in Section 3.2.1.

Due to parity invariance of the strong interactions, $t\bar{t}$ pairs are produced almost unpolarised at the LHC. Consequently, the individual top and antitop spin polarisations, B_i^+ and B_i^-, vanish at LO. Nevertheless, since the top quark has a lifetime shorter than the typical spin decorrelation timescale, the spin information is preserved until decay, and top-quark pairs are predominantly produced with correlated spins. In particular, the production mechanisms at the LHC tend to favour configurations where both the top quark and the antitop quark have the same-helicity, either both left-handed ($\downarrow\downarrow$) or both right-handed ($\uparrow\uparrow$)[2]. This is nicely illustrated in Figure 3.4 taken from Ref. [69]. At higher $m_{t\bar{t}}$, the top quarks become more relativistic, and opposite-helicity configurations ($\uparrow\downarrow$ or $\downarrow\uparrow$) become increasingly significant. This preference for aligned helicities leads to an asymmetry between the possible helicity combinations of the top-quark pair, characterised by spin alignment ($\uparrow\uparrow$ or $\downarrow\downarrow$) and spin anti-alignment ($\uparrow\downarrow$ or $\downarrow\uparrow$). These spin correlation patterns are encoded in the spin correlation matrix $\mathbf{C}$.

As seen from Equation (3.4), the complete spin density matrix of the $t\bar{t}$ system can be reconstructed by analysing the angular distributions of the top-quark decay products. This requires identifying which of the top-quark decay products exhibit angular distributions most strongly correlated with the spin direction of their parent quark, and thereby carry the most spin information. The decay product

[1] Polarisation refers to a non-random alignment of the spin relative to some reference direction or axis. A particle is considered polarised when a specific direction of spin is preferred over others.

[2] Helicity, defined as the projection of the spin of a particle onto the direction of its momentum, characterises whether the spin is aligned or anti-aligned with the direction of motion of the particle in a given frame. Particles with spin aligned along their momentum have right-handed helicity ($\uparrow$), while those with spin anti-aligned have left-handed helicity ($\downarrow$). Note that for particles with mass, helicity is frame-dependent.

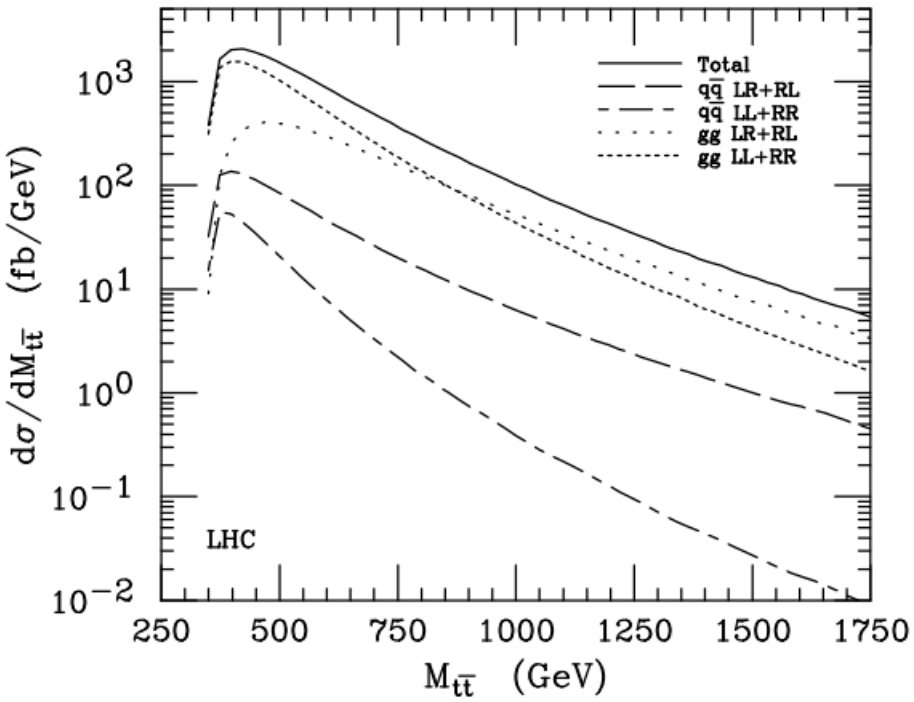


Figure 3.4: Differential cross-section for $t\bar{t}$ production as a function of the $t\bar{t}$ invariant mass at the LHC with $\sqrt{s}$ = 14 TeV, separated by left-handed and right-handed helicity configurations for both production modes [69].

that retains the highest sensitivity to the spin orientation of the parent top quark is known as the best ***spin analyser***. The following section outlines the available choices of spin analysers for both the semi-leptonically and hadronically decaying top quarks in the lepton+jets decay channel.

3.2.1 Top-quark spin analysers

Since the angular distributions of top-quark decay products carry information of the parent quark spin, selecting an optimal spin analyser is essential to achieve optimal sensitivity in quantum measurements using $t\bar{t}$. The extent to which a given decay product retains the spin information of the parent top quark is quantified by its ***spin analysing power***, denoted by α (e.g., α_1 and α_2 in Equation (3.4)). This parameter measures the strength of the correlation between the emission direction of the decay product and the spin direction of the decaying top quark. The value of α ranges from -1 to 1, where $\alpha = \pm 1$ corresponds to a perfect spin analyser and $\alpha = 0$ indicates no spin sensitivity. For antiparticles, the sign of α is reversed.

A convenient way to compute α for top-quark decay products is by considering the differential decay distribution for an ensemble of top quarks at rest with a polarisation vector $\vec{P}$, where $0 \leq |\vec{P}| \leq 1$ [70]:

$$\frac{1}{\Gamma}\frac{d\Gamma}{d\cos\theta} = \frac{1}{2}\left(1 + |\vec{P}|\alpha_p \cos\theta\right). \tag{3.5}$$

Here, θ is the angle between $\vec{P}$ and the direction $\hat{p}$ of a given decay product, Γ is the partial width for the corresponding decay of unpolarised top quarks and α_p is the spin analysing power of the decay

product. Correspondingly the expectation value:

$$\langle \cos\theta \rangle = \langle \hat{p} \cdot \hat{P} \rangle = \frac{1}{\Gamma} \int d\Gamma \cos\theta = \frac{\alpha_p |\vec{P}|}{3},$$

provides a convenient way to extract α_p given a known $|\vec{P}|$. Theoretically, the differential rate $\frac{d\Gamma}{d\cos\theta}$ is derived from the squared matrix element $|\mathcal{M}^2|$, which encodes the spin structure of the decay. A dedicated calculation of the matrix element and phase-space integration, using the angular dependence from the matrix element, is required to determine α_p for each decay product. More details on the calculations can be found in Ref. [70].

Spin analysers for semi-leptonically decaying top quarks While multiple decay products emerge from top-quark decays via the weak interaction, they differ in their ability to reflect the spin orientation of the parent top quark. This distinction is primarily governed by the vector minus axial-vector (V–A) structure of the charged-current weak interaction, which mediates the decay $t \to bW^+$ (or $\bar{t} \to \bar{b}W^-$). The V–A nature of the interaction means that the weak force couples only to left-handed fermions and right-handed antifermions[3].

The W boson produced in this decay is a massive spin-1 particle and therefore has three possible polarisation or spin states: $+1$, 0, and -1. In contrast, the b-quark, being a spin-$\frac{1}{2}$ fermion, has two spin states: $+\frac{1}{2}$ and $-\frac{1}{2}$. Since the weak interaction couples only to left-handed b-quarks, the helicity of the b-quark is effectively fixed to be left-handed.

As a result, in the top-quark rest frame, the decay $t \to bW^+$ can proceed through three possible W-boson helicity configurations: right-handed, left-handed, and longitudinal (zero) helicity. These configurations are illustrated in Figure 3.5. Due to the V–A structure of the weak interaction and angular momentum conservation, not all helicity states are equally likely. In particular, the decay into a right-handed W boson is highly suppressed. This happens because such a configuration would then require the b-quark to have right-handed helicity if it were massless, which is in conflict with angular momentum conservation.

Therefore, the decay predominantly produces W bosons with left-handed or longitudinal helicity. As seen in Figure 3.5, the spin of the W boson is strongly correlated with that of the parent top quark (the double arrows showing W boson and top quark spin directions). However, since the W boson can be produced in multiple helicity states, its emission direction does not fully preserve the spin information of the top quark. This leads to a dilution of the spin correlation, resulting in a spin analying power of $\alpha_W \approx 0.41$.

The b-quark, on the other hand, exhibits spin alignment with the top quark depending on the W-boson helicity: their spins are parallel for longitudinal W bosons and anti-parallel for left- and right-handed ones. Furthermore, the direction of motion of the b-quark is anti-parallel to its spin. As a result, it offers only modest spin sensitivity like the W boson, with a spin analysing power of $\alpha_b \approx -0.41$ [70].

In leptonic decays of the W boson ($W^+ \to l^+ \nu_l$), the helicity state of the charged lepton is fixed due

[3] Strictly speaking, the weak interaction couples to fermions of definite chirality rather than helicity. However, for highly relativistic particles, where the particle energy is much larger than its mass, the chirality and helicity states approximately coincide. In the context considered here, helicity therefore provides an adequate description of the relevant spin configurations.

to V–A interaction, allowing it to fully inherit the spin information of the W boson, and hence the top quark, with minimal dilution. Moreover, since it does not undergo hadronisation and can be more easily detected in contrast to the neutrino that escapes detection, it serves as the most effective spin analyser, with a spin analysing power close to unity ($\alpha_\ell \approx 1$) [70].

Therefore, in the case of the semi-leptonically decaying top quark, the ***charged lepton*** is used as the optimal spin analyser in this measurement.

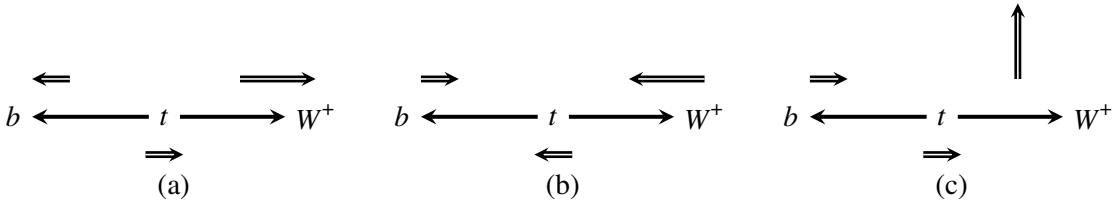


Figure 3.5: Illustration of a top-quark decay, in the top-quark rest frame, into a b quark and a W^+ boson with (a) right-handed helicity (b) left-handed helicity and (c) longitudinal (zero) helicity [71, 72]. The single arrows indicate the directions of the particle momenta, while the double arrows denote the spin directions of the top quark (lower arrow) and of the b-quark and W boson (upper arrows).

Spin analysers for hadronically decaying top quarks If the W boson, coming from the top quark, decays hadronically $t \to b\bar{d}u$, where $d(u)$ generically refers to d, s (u, c) quarks, the ***down-type quark*** $\bar{d}$ is the parton that corresponds to the charged lepton in the hadronic decay. Therefore, using the same argument as for the lepton above, at LO the down-type quark carries a spin analysing power of $\alpha_{\bar{d}} \approx 1$. The up-type quark, on the other hand, carries a much lower spin-analysing power of $\alpha_u \approx -0.31$ at LO [70]. Since the final state ultimately involves two jets coming from the hadronic W-boson decay, it is necessary to correctly identify at the detector level which jet originates from the down-type quark. The identification of a jet as originating from a specific quark flavour is referred to as *tagging*. This is a very important procedure used in various measurements involving top quarks, such as the one presented in this thesis. Dedicated algorithms exist to identify jets originating from b-quarks, commonly referred to as b-tagging algorithms (Section 6.4.6). These algorithms exploit the distinctive properties of b-jets to distinguish them from other quark flavours. However, no dedicated tools currently exist for directly tagging down-type jets. Instead, the identification is done indirectly by tagging the accompanying jet coming from a charm quark (c-jet) using the available taggers, with the other jet being assumed to be the down-type jet. Due to inefficiencies in c-jet tagging (Section 6.4.6) and the fact that the W boson decays to a c-jet approximately 50% of the time, not all events involving this analyser can be successfully captured. As a result, alternative spin analyser approaches can be considered for hadronic top quarks. One such option is using the ***b-jet*** coming from the top-quark decay, which can be reliably identified via b-tagging techniques, although it offers a lower spin analysing power ($\alpha_b \approx -0.41$). Another possibility is to use the least energetic jet, called the ***soft jet***, from the W-boson decay, which at LO has a spin analying power of $\alpha_s \approx -0.51$ [70].

Another method involves combining information of both jets from the W-boson decay to construct a spin analyser known as the ***optimal hadronic polarimeter*** [73]. This approach exploits the polarised nature of the W boson, enabling the definition of the polar decay angle within the W-boson rest frame, commonly referred to as the helicity angle, $\theta_{W\,\mathrm{hel}}$. The angle tends to exhibit a bias due to the polarisation states of the W boson. By defining the z-axis of the W-boson decay to point opposite to

the direction of the b quark (i.e, $\hat{z} \equiv -\hat{b}$) in the W-boson rest frame, the distribution of the helicity angle can be written as:

$$\rho(c_{W\text{hel}}) \equiv \frac{3}{8} f_R (1 + c_{W\text{hel}})^2 + \frac{3}{4} f_0 (1 - c_{W\text{hel}}^2) + \frac{3}{8} f_L (1 - c_{W\text{hel}})^2,$$

where $c_{W\text{hel}}$ is the cosine of the helicity angle. In this case, the angle is defined as the direction of the down-type quark with respect to $-\hat{b}$, as indicated in Figure 3.6. The fractions of right-handed, zero, and left-handed W bosons in the top frame are denoted as f_R, f_0, and f_L respectively. As briefly mentioned before, f_R is nearly zero by the argument in Section 3.2.1, and:

$$f_0 \simeq \frac{m_t^2}{m_t^2 + 2m_W^2} \simeq 0.70,$$
$$f_L \simeq \frac{2m_W^2}{m_t^2 + 2m_W^2} \simeq 0.30.$$

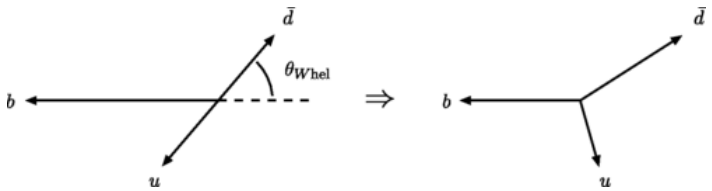


Figure 3.6: The top-quark decay system as viewed in the W-boson rest frame (left), and boosted back into the top-quark rest frame (right). The helicity angle, $\theta_{W\text{hel}}$, is defined as the direction of the down-type quark with respect to $-\hat{b}$ [73].

Given the challenges of identifying the jets from the W-boson decay as originating from a d- or u-type quark, a more practical labelling scheme is used based on their kinematic properties after boosting from the W rest frame to the top-quark rest frame. The quark emitted forward in the W frame becomes more energetic (the "hard" jet) and is typically more angularly separated from the accompanying b-quark, while the backward-emitted quark becomes less energetic (the "soft" jet) and more closely aligned with the b-quark. Although QCD showering slightly distorts this picture, the distinction between hard and soft jets remains useful. Since the jets are no longer labelled by flavour, the d-quark is treated as having a certain probability of ending up in either the soft or hard jet:

$$p(d \to q_{\text{soft}}) = \frac{\rho(-|c_{W\text{hel}}|)}{\rho(|c_{W\text{hel}}|) + \rho(-|c_{W\text{hel}}|)},$$
$$p(d \to q_{\text{hard}}) = \frac{\rho(|c_{W\text{hel}}|)}{\rho(|c_{W\text{hel}}|) + \rho(-|c_{W\text{hel}}|)}.$$

Taking the soft-quark direction as $\hat{q}_{\text{soft}}$ and the hard-quark direction as $\hat{q}_{\text{hard}}$, the optimal hadronic polarimeter direction can be written as a weighted combination of these two:

$$\vec{q}_{\text{opt}}(|c_{W\text{hel}}|) \equiv p(d \to q_{\text{soft}})\hat{q}_{\text{soft}} + p(d \to q_{\text{hard}})\hat{q}_{\text{hard}},$$

and the spin anlysing power for a given value of $|c_{W\,\mathrm{hel}}|$ is the length of this vector:

$$\alpha_{\mathrm{opt}}(|c_{W\,\mathrm{hel}}) = |\vec{q}_{\mathrm{opt}}(|c_{W\,\mathrm{hel}}|)|.$$

Figure 3.7 shows the spin analysing power of this optimal hadronic polarimeter as a function of $|c_{W\,\mathrm{hel}}|$, along with several other hadronic spin analyser choices. It can been seen that the optimal polarimeter maintains an approximately constant spin analysing power of $\alpha_{\mathrm{opt}} \approx -0.64$, making it a strong candidate for analysing the spin information of the hadronic top quark.

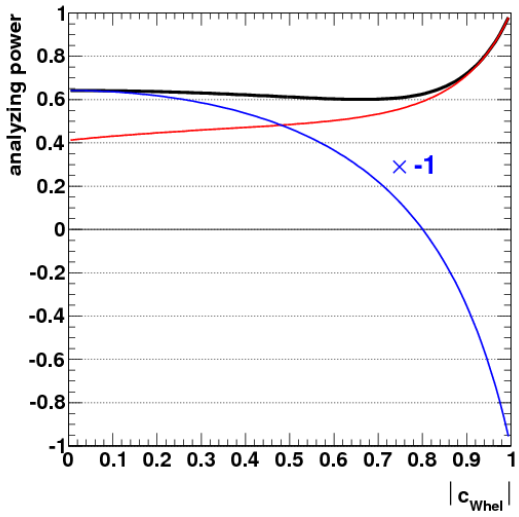


Figure 3.7: Spin analysing powers as a function of $|c_{W\,\mathrm{hel}}|$ for the optimal hadronic polarimeter q_{opt} (black), soft quark q_{soft} (red), and b-quark (blue) [73].

Another potential spin analyser is the reconstructed ***hadronic W boson*** from the top quark decay ($\alpha_W \approx 0.41$), which is particularly valuable in topologies where the two light jets from the W-boson decay are nearly indistinguishable. Table 3.2 summarises the various spin analyser options and their corresponding spin analysing powers at LO and NLO. At NLO, additional radiation and virtual corrections can dilute the correlation between the top-quark spin direction and the emission direction of a given spin analyser, slightly reducing the spin analysing powers. For quark-based spin analysers, the NLO values corresponding to their jets, calculated in Ref. [70], are also included. These are lower than the spin analysing powers of the corresponding quarks, which is expected since some correlation is lost when quarks hadronise to form jets. The down-type jets show a larger reduction in spin analysing power relative to their corresponding partons, whereas the b-jets are affected only slightly. This could arise because the spin correlation of the down-type quark with its parent top quark is nearly maximal, so even small distortions in the W-decay geometry due to NLO radiation or jet clustering can significantly dilute the correlation. In contrast, the b-quark already has a moderate correlation leading to only a minor change in its spin analysing power.

	b-quark	W^+	l^+	soft quark	down-type quark ($\bar{d}$)	up-type quark (u)
α_i(LO)	-0.41	0.41	1	0.51	1	-0.31
α_i(NLO)	-0.392	0.392	0.998	-	0.966	-0.317
α_i(NLO) jets	-0.391	0.391	-	-	0.938	-0.303

Table 3.2: Spin analysing powers of potential top-quark spin analysers [70, 74].

3.2.2 Spin correlations in top-quark pairs

For the $t\bar{t}$ system, the analogous form of the spin density matrix in Equation (3.3) can be written as the so-called production spin density matrix $R(m_{t\bar{t}}, \hat{k})$ [75]:

$$R = \tilde{A}\mathbb{I}_4 + \sum_{i=1}^{3}\left(\tilde{B}_i^+\sigma^i \otimes \mathbb{I}_2 + \tilde{B}_i^-\mathbb{I}_2 \otimes \sigma^i\right) + \sum_{i,j=1}^{3}\tilde{C}_{ij}\sigma^i \otimes \sigma^j. \tag{3.6}$$

Here, $\tilde{A}$ is an additional parameter that characterises this matrix and ensures proper normalisation of the differential cross-section for $t\bar{t}$ production at fixed centre-of-mass energy and top-quark direction. The spin density matrix characterises the spin correlations between the top and antitop quarks at production and depends on the invariant mass of the $t\bar{t}$ system, $m_{t\bar{t}}$ and the top-quark direction, denoted by $\hat{k}$ in the $t\bar{t}$ centre-of-mass (CM) frame. More precisely, the matrix elements depend only on two dimensionless variables: the velocity β of the top quark, expressed in natural units (where speed of light, $c = 1$), which is determined by $m_{t\bar{t}}$ via $\beta = \sqrt{1 - \frac{4m_t^2}{m_{t\bar{t}}^2}}$ and the cosine of the production angle, $\cos\Theta = \hat{k} \cdot \hat{p}$, with $\hat{p}$ being the unitary vector along a chosen proton-beam direction.

The degree of measured spin correlation in $t\bar{t}$ production depends on the choice of the spin quantisation basis. Using an optimal basis increases the separation between correlated and uncorrelated $t\bar{t}$ pairs, thus making the correlation C maximal. Several spin bases can be defined, with two of the most commonly used illustrated in Figure 3.8.

In the ***beam basis*** (shown on the right), the spin quantisation axes of the top and antitop quarks are aligned with the directions of the incoming partons. This basis is particularly effective at lower $m_{t\bar{t}}$ values, where the production process $gg \to t\bar{t}$ with same-helicity $t\bar{t}$ states tends to dominate (see Figure 3.4).

On the other hand, the ***helicity basis***, denoted by the orthonormal set $\hat{k}, \hat{n}, \hat{r}$ (left of Figure 3.8), is defined in the $t\bar{t}$ CM frame. Here, $\hat{r} = (\hat{p} - \cos\Theta\hat{k})/\sin\Theta$ and $\hat{n} = \hat{r} \times \hat{k}$. The antitop spin axes are defined analogously. This basis is particularly suited for measuring the individual components of the spin correlation matrix C, especially in the high $m_{t\bar{t}}$ region where opposite-helicity $t\bar{t}$ states dominate.

In the SM, due to CP invariance of the strong interaction, the spin correlation matrix $\tilde{C}_{ij}$ is symmetric, and the polarisation vectors satisfy $\tilde{B}_i^+ = \tilde{B}_i^-$. At LO, since QCD is parity- and CP-conserving, the net polarisations vanish, $\tilde{B}_i^\pm = 0$, and spin components along the n-axis are uncorrelated with those in the transverse directions, $\tilde{C}_{nr} = \tilde{C}_{nk} = 0$. A parity-invariant theory like LO QCD cannot have observables that change sign under a parity transformation, and longitudinal-transverse spin correlations are parity-odd. Consequently, only five parameters are needed to characterise the spin density matrix at LO: $\tilde{A}, \tilde{C}_{kk}, \tilde{C}_{nn}, \tilde{C}_{rr}, \tilde{C}_{kr}$.

Spin correlations in $t\bar{t}$ can be measured through several methods, the most natural being via the

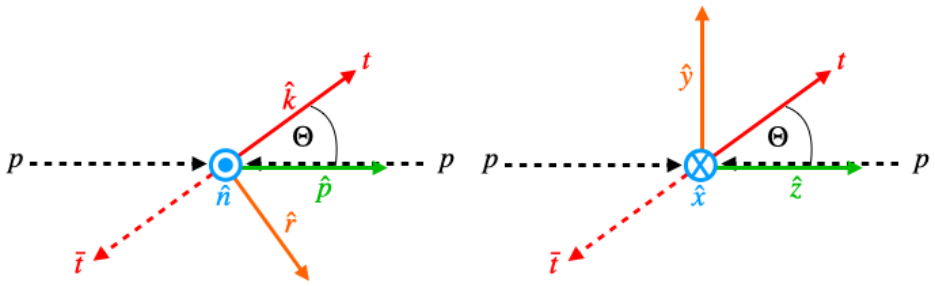


Figure 3.8: Two different orthonormal bases in the CM frame: helicity basis on the left and beam basis on the right [75].

study of angular distributions of the decay products. A specific component C_{ij} of the spin correlation matrix can be extracted from the distribution of the product $\cos\theta_i^a \cos\theta_j^b$, where $\cos\theta_i^a$ denotes the cosine of the angle between the the spin analyser a and the spin quantisation axis i. The ATLAS collaboration measured all independent components of the $t\bar{t}$ spin density matrix at $\sqrt{s} = 8$ TeV in the dilepton channel [76], representing the only existing measurement by ATLAS in this channel. The CMS collaboration also performed the corresponding measurement at $\sqrt{s} = 13$ TeV [77]. Another approach to quantify spin correlations involves measuring the fraction of SM-like spin correlation, f_{SM}, through the azimuthal angular separation $\Delta\phi$ between the two top quark spin analysers. The most recent ATLAS result [78], at $\sqrt{s} = 13$ TeV in the dilepton channel, employs the angular variable $\Delta\phi(\ell^+, \ell^-)$, the azimuthal separation between the charged leptons from the top and antitop decays in the lab frame, as a key observable. CMS has also performed a similar measurement [77]. A combination of these two measurements [79] provides the latest determination of azimuthal spin correlations. The only existing complete spin density matrix measurement in the lepton+jets channel by ATLAS was performed at $\sqrt{s} = 7$ TeV [80]. CMS performed the corresponding measurement at $\sqrt{s} = 13$ TeV [24], extracting spin density matrix components both inclusively and as a function of $m_{t\bar{t}}$, the top-quark transverse momentum $p_T(t)$, and $|\cos\Theta|$. All of these measurements are consistent with the SM predictions.

These spin correlation measurements not only test the predictions of the SM but also reflect the inherently quantum nature of the $t\bar{t}$ system. In particular, they can point to the presence of correlations between the top and antitop spins that are stronger than classically expected, correlations that are naturally described in the framework of quantum entanglement.

3.2.3 Quantum entanglement in top-quark pairs

Referring back to Equation (3.2), if the subsystems of the state cannot be expressed as a separable product, then the state is termed entangled. Such states exhibit non-classical quantum correlations, for which the quantum state of the full system cannot be factorised into independent states of the individual particles. This phenomenon is known as ***quantum entanglement***.

In bipartite systems, a key criterion for determining separability is the ***Peres-Horodecki criterion*** [81, 82], which states that if a state ρ is separable, then the partial transpose of ρ (e.g with respect to the

second subsystem):

$$\rho^{\mathrm{T}_b} = \sum_{n=1}^{N} p_n \rho_n^a \otimes (\rho_n^b)^{\mathrm{T}},$$

must also represent a physical state, i.e., a non-negative operator with unit trace. Thus, if ρ^{T_b} possesses negative eigenvalues, the state is entangled. This means that subsystem a cannot be described independently of subsystem b.

A quantitative measure of entanglement is provided by the concurrence $C[\rho]$, which relates to the degree of entanglement. It is defined as [83]:

$$C[\rho] \equiv \max\left(0, \lambda_1 - \lambda_2 - \lambda_3 - \lambda_4\right),$$

where λ_i are the eigenvalues, in decreasing order, of the matrix $C(\rho) = \sqrt{\sqrt{\rho}\tilde{\rho}\sqrt{\rho}}$ with $\tilde{\rho} = (\sigma_2 \otimes \sigma_2)\rho^*(\sigma_2 \otimes \sigma_2)$, and ρ^* the complex conjugate of the density matrix, ρ. The concurrence takes values between 0 and 1 for entangled states, with 1 indicating maximal entanglement.

Applying the Peres-Horodecki criterion to the production spin density matrix in Equation (3.6), after considering $\tilde{B}_i^{\pm} = 0$ and $\tilde{C}_{ij} = \tilde{C}_{ji}$, yields a set of necessary conditions for a state to be entangled [75]:

$$|C_{rr} + C_{kk}| - C_{nn} > 1,$$
$$|C_{rr} - C_{kk}| + C_{nn} > 1.$$

These conditions can be related to the concurrence as

$$C[\rho] = \begin{cases} \frac{1}{2}\max\left(|C_{rr} + C_{kk}| - 1 - C_{nn},\, 0\right) & \text{if } C_{nn} \leq 0. \\ \frac{1}{2}\max\left(|C_{rr} - C_{kk}| - 1 + C_{nn},\, 0\right) & \text{if } C_{nn} \geq 0. \end{cases} \tag{3.7}$$

Depending on the $t\bar{t}$ production mechanism, boundaries on the concurrence can be established to distinguish separable and entangled regions of phase space. Figure 3.9(a) illustrates the concurrence for the gluon fusion mode $\rho^{gg}(m_{t\bar{t}}, \hat{k})$. The black lines mark the transitions between separable and entangled regions, with entanglement present in the lower-left and upper-right corners.

This behaviour can be attributed to how the spin polarisations of the gluon pair align in different directions across phase space. At the production threshold ($m_{t\bar{t}} = 2m_t$), where the relative velocity of the $t\bar{t}$ pair is small (lower-left region of Figure 3.9(a)), the system forms a ***spin-singlet*** state given by:

$$\rho^{gg}(2m_t, \hat{k}) = |\Psi_0\rangle\langle\Psi_0|, \quad |\Psi_0\rangle = \frac{1}{\sqrt{2}}\left(|\uparrow_{\hat{n}}\downarrow_{\hat{n}}\rangle - |\downarrow_{\hat{n}}\uparrow_{\hat{n}}\rangle\right),$$

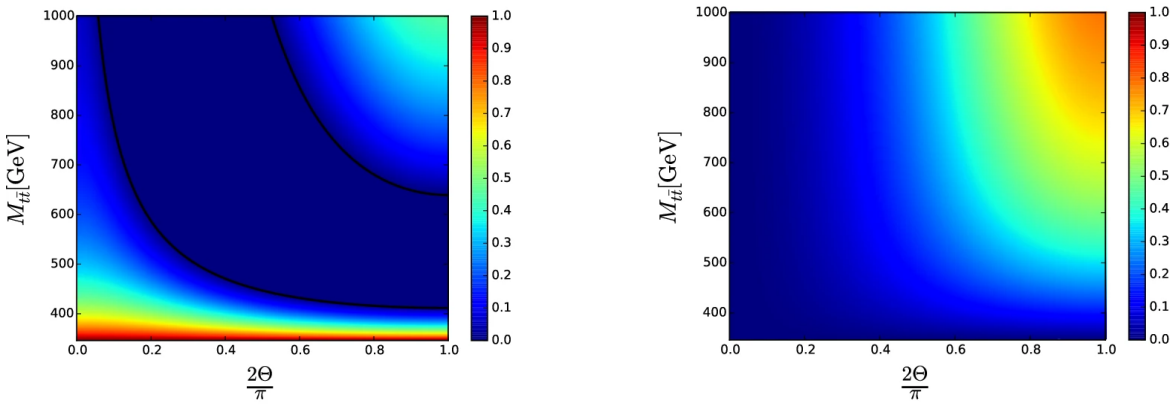
where $\downarrow_{\hat{n}}\uparrow_{\hat{n}}$ denote the spin eigenstates along the direction $\hat{n}$. This antisymmetric spin configuration is maximally entangled, since the two top quark substates cannot be factorised individually.

At very high $m_{t\bar{t}}$ and large production angles (upper-right region of Figure 3.9(a)), the $t\bar{t}$ pair is again maximally entangled due to the formation of spin-triplet pure states with symmetric spins:

$$\rho^{gg}(\infty, \hat{n} \times \hat{p}) = |\Psi_\infty\rangle\langle\Psi_\infty|, \quad |\Psi_\infty\rangle = \frac{1}{\sqrt{2}}\left(|\uparrow_{\hat{n}}\downarrow_{\hat{n}}\rangle + |\downarrow_{\hat{n}}\uparrow_{\hat{n}}\rangle\right).$$

Between these two regions, the top quarks are predominantly produced in non-entangled states.

In contrast, for $q\bar{q}$-initiated production, the state remains entangled across most of the phase space.



(a) $gg \to t\bar{t}$ (b) $q\bar{q} \to t\bar{t}$

Figure 3.9: Measure of entanglement, termed concurrence, as a function of the invariant mass of the $t\bar{t}$ system, $m_{t\bar{t}}$ (denoted by $M_{t\bar{t}}$ in these figures), and the production angle Θ in the $t\bar{t}$ CM frame. Black lines represent the boundaries between separability and entanglement [75].

However, at $\Theta = 0$ or near the production threshold, the entanglement becomes suppressed, and the $t\bar{t}$ pairs emerge with spins aligned along the beam axis in correlated but separable mixed states:

$$\rho^{q\bar{q}}(2m_t, \hat{k}) = \rho^{q\bar{q}}(m_{t\bar{t}}, \hat{p}) = \frac{1}{2}\left(|\uparrow_{\hat{p}}\uparrow_{\hat{p}}\rangle\langle\uparrow_{\hat{p}}\uparrow_{\hat{p}}| + |\downarrow_{\hat{p}}\downarrow_{\hat{p}}\rangle\langle\downarrow_{\hat{p}}\downarrow_{\hat{p}}|\right).$$

Thus, in these limits, the degree of entanglement is expected to be small for $q\bar{q}$ production, as illustrated in Figure 3.9(b). In the opposite limit of very high energies and large production angles, the system again approaches a maximally entangled state, similar to the case of gg fusion. In this region, however, the degree of entanglement is larger for $q\bar{q}$-initiated production than for gg fusion, indicating that in this region the maximally entangled states are predominantly produced through the $q\bar{q}$ channel.

Based on these observations, two distinct regions of phase space emerge in which entanglement can be effectively probed: the region near the production threshold, dominated by gg fusion and referred to as the ***threshold region***, and the region at high $t\bar{t}$ invariant masses and large production angles, dominated by $q\bar{q}$ production and referred to as the ***boosted region***. The term "boosted" reflects that, in this region, the top quarks move very fast in the laboratory frame with large transverse momenta p_T and carry large Lorentz boosts ($\gamma = E_t/m_t \gg 1$).

In the threshold region, the spin correlation matrix components satisfy $C_{rr} < 0$, $C_{kk} < 0$, and $C_{nn} < 0$. Using these in Equation (3.7) yields the entanglement condition:

$$-C_{rr} - C_{kk} - C_{nn} - 1 > 0. \tag{3.8}$$

In contrast, in the boosted region: $C_{rr} > 0$, $C_{kk} > 0$, and $C_{nn} < 0$. This leads to a different form of the inequality:

$$C_{rr} + C_{kk} - C_{nn} - 1 > 0. \tag{3.9}$$

These two criteria can be recast in terms of two measurable observables [84]:

$$D = \frac{1}{3}(C_{rr} + C_{kk} + C_{nn}),$$
$$D_3 = \frac{1}{3}(-C_{rr} - C_{kk} + C_{nn}).$$

Both D and D_3 can be extracted from the differential cross-section characterising the angular separation between the spin analysers of the two top quarks:

$$\frac{1}{\sigma}\frac{d\sigma}{d\cos\varphi} = \frac{1}{2}(1 + \alpha_1\alpha_2 D\cos\varphi),$$
$$\frac{1}{\sigma}\frac{d\sigma}{d\cos\theta'} = \frac{1}{2}(1 + \alpha_1\alpha_2 D_3\cos\theta').$$

Here, φ is the angle between the spin-analyser directions in the respective rest frames of the parent top and antitop quarks, after boosting to the $t\bar{t}$ rest frame. The angle θ' is defined analogously in the helicity basis, and corresponds to the angle between one spin-analyser direction and the reflected direction of the other analyser in the basis defining D_3 [84]. The parameters α_1 and α_2 denote the spin analysing powers of the respective analysers.

The observables can be extracted by taking the mean[4] of the corresponding angular distributions:

$$D = 3 \cdot \frac{\langle\cos\varphi\rangle}{\alpha_1\alpha_2}, \tag{3.10}$$
$$D_3 = 3 \cdot \frac{\langle\cos\theta'\rangle}{\alpha_1\alpha_2}. \tag{3.11}$$

The ***entanglement criterion***, from Equations (3.8) and (3.9), require that the corresponding observable satisfies:

$$D_{\text{eff}} < -\frac{1}{3}. \tag{3.12}$$

where D_{eff} refers to either D (in the threshold region) or D_3 (in the boosted region).

From these observables, the concurrence can be directly obtained as $C[\rho] = \frac{1}{2}\max(-3D - 1, 0)$ in the threshold region and $C[\rho] = \frac{1}{2}\max(-3D_3 - 1, 0)$ in the boosted region. Since D is proportional to the trace of the spin correlation matrix C, it is basis-independent and can be measured in either the beam or helicity basis. D_3, on the other hand, is basis-dependent. In the boosted region entanglement measurement presented in this thesis, the helicity basis is used.

In summary, the threshold region, dominated by gg fusion at the LHC, becomes particularly useful for accessing entanglement under experimentally favourable conditions. Entanglement in this region in top-quark pairs was measured for the first time at the LHC in the dilepton channel in ATLAS [17]. This was followed by results from CMS in both the dilepton [23] and lepton+jets channel [24]. CMS has also performed entanglement measurement using the lepton+jets channel in the boosted regime [24], which in contrast offers a complementary window into entanglement in the extreme high-energy limit of phase space where the top quarks are space-like separated. This means that the two top quarks are

[4] In principle, the observables can also be extracted from the slope of the corresponding distributions, which would be optimal for linear distributions. However, in experimental measurements the distributions deviate from linearity due to detector and reconstruction effects, making this method less precise than the mean-based approach (Section 7.7).

separated in such a way that no signal travelling at or below the speed of light could connect them. While this regime is experimentally more challenging, probing entanglement and performing quantum state tomography not only extend the reach of such studies, but would also set the stage for testing other quantum effects such as measuring Bell correlated states which become accessible only in this space-like separated regime.

3.2.4 Measuring Bell correlated states with top quarks

While entanglement is a fundamental property of QM, entangled particles can exhibit correlations that cannot be explained by any theory that extends the classical picture, by introducing hidden variables that explain these stronger correlations in a classical manner. This tension between QM and classical notions was first highlighted by Einstein, Podolsky, and Rosen (EPR) in their famous 1935 paradox [13], which argued that QM must be incomplete if it allows instantaneous effects between distant particles, a phenomenon which Einstein famously referred to as "spooky action at a distance".

The paradox showed that QM predicts perfect correlations between measurements on two spatially separated particles in an entangled state, which is in conflict with the classical principles of locality. It proposed the existence of local hidden variables, which are unknown parameters that would restore determinism and locality.

In 1964, John Bell formalised this paradox into a testable criterion, now known as ***Bell's inequality*** [12]. Any theory based on local hidden variables must satisfy this inequality, whereas certain quantum correlations prediced by QM can violate it. For a bipartite system of two qubits, the most well-known form of Bell's inequality is the ***Clauser-Horne-Shimony-Holt (CHSH) inequality***, which places bounds on correlations between measurements on subsystems A and B [85]:

$$\langle A_1 B_1 \rangle - \langle A_1 B_2 \rangle + \langle A_2 B_1 \rangle + \langle A_2 B_2 \rangle \leq 2 \tag{3.13}$$

Here, $\langle A_i B_j \rangle$ denotes the expectation value of simultaneous measurements A_i on subsystem A and B_i on subsystem B. A bipartite quantum state that violates this condition exhibits Bell non-locality, meaning that the observed correlations between the two subsystems cannot be explained by any local hidden variable theory.

For the $t\bar{t}$ system, the inequality given in Equation 3.13 can be written in the helicity basis, where it corresponds to the following relation between the spin correlation coefficients [84]:

$$B = C_{nn} - C_{rr}.$$

Bell inequality violation occurs if $B > \sqrt{2}$. Consider the azimuthal angles ϕ_1 and ϕ_2 of the two spin analysers. In the helicity basis, the azimuthal angle ϕ is defined around the $\hat{k}$ direction, with $\phi = 0$ lying in the $\hat{n} - \hat{k}$ plane. Defining the combinations

$$\phi_+ = \frac{\phi_1 + \phi_2}{2}, \quad \phi_- = \frac{\phi_1 - \phi_2}{2},$$

in Equation (3.4), the differential cross-section becomes:

$$\frac{1}{\sigma}\frac{d\sigma}{d\phi_+ d\phi_-} = \frac{1}{2\pi^2} + \frac{\alpha_1\alpha_2}{32}\left(\frac{C_{nn}+C_{rr}}{2}\cos(2\phi_-) + \frac{C_{nn}-C_{rr}}{2}\cos(2\phi_+) + \frac{C_{nn}+C_{rn}}{2}\sin(2\phi_+) + \frac{C_{rn}-C_{nr}}{2}\sin(2\phi_-)\right).$$

The term involving $C_{nn} - C_{rr}$ is the only one that is an even function of ϕ_+, which enables its extraction through the asymmetry:

$$B = C_{nn} - C_{rr} = \frac{16}{\pi\alpha_1\alpha_2}(A_{\cos(2\phi_+)}) \tag{3.14}$$

where the asymmetry $A_{\cos(2\phi_+)}$ is defined as:

$$A_{\cos(2\phi_+)} = \frac{N(\cos(2\phi_+) > 0) - N(\cos(2\phi_+) < 0)}{N(\cos(2\phi_+) > 0) + N(\cos(2\phi_+) < 0)}$$

Here, $N(\cos(2\phi_+) > 0)$ and $N(\cos(2\phi_+) < 0)$ denote the number of events where $\cos(2\phi_+)$ is positive or negative, respectively. Hence, the parameter B can also be extracted using a single variable, $\cos(2\phi_+)$.

In summary, a violation of the Bell inequality would indicate quantum entanglement and non-local spin correlations which cannot be explained by any classical or local hidden variable theories. The converse, on the other hand, is not always true: the presence of spin correlations or even entanglement does not necessarily imply a Bell inequality violation. These effects are related through a hierarchical structure:

$$\text{Bell inequality violation} \subseteq \text{Entanglement} \subseteq \text{Spin correlations}$$

This means that every Bell-inequality violating system is entangled and exhibits spin correlations, but not all entangled or spin-correlated systems will violate the Bell inequality. This happens because some entangled states do not produce sufficiently strong correlations to exceed the classical bound. Similarly, spin correlations can also exist in classical or separable systems and are therefore not always quantum entangled.

In the context of top-quark pairs, this hierarchy becomes even more sensitive to the phase space under consideration. The strength of observed spin correlations depends heavily on the kinematic region that is probed, the selections applied, and the subsequent detector effects.

While Bell inequality violation provides a powerful test of non-classical correlations, performing such a measurement with top quarks at the LHC is not possible in a loophole-free manner. While top quarks can be produced with very high $m_{t\bar{t}}$, an environment ideal to measure Bell inequality violation, a direct Bell test at colliders is currently not feasible. In collider experiments, the spin polarisation of top quarks cannot be determined on an event-by-event basis. The spin correlations can only be accessed indirectly from the angular distributions of the decay products. Refs. [86–88] argued that Bell tests cannot be performed at colliders, since the experimentally accessible observables (such as particle momenta) are commuting, whereas violation requires measurements of non-commuting quantities. Therefore, ***the condition $B > \sqrt{2}$ in $t\bar{t}$ systems cannot be established in a loop-hole free manner at***

colliders. Nevertheless, one can still probe a Bell correlated state, i.e. a quantum state that would, in principle, violate Bell inequality if a proper Bell test were possible in collider experiments, even though the violation itself cannot be tested at a collider. Measuring the expression in Equation 3.14 can therefore provide an understanding of Bell-correlated states and also help scrutinise the SM in the high-$m_{t\bar{t}}$ region.

Unlike Bell inequality violation, however, entanglement itself remains experimentally accessible under the assumptions of SM, allowing one to investigate whether the produced $t\bar{t}$ states are entangled in a given region of phase space without a Bell test.

3.3 Exploring quasi-bound-state effects in the threshold region

Entanglement in the $t\bar{t}$ system was first measured by ATLAS in the dilepton channel near the threshold region [17]. While this measurement marked a significant milestone, discrepancies were observed between the expected and measured values of the entanglement observable D, suggesting potential gaps in the theoretical understanding of the QCD dynamics that govern the low-$m_{t\bar{t}}$ regime. Among the various theoretical possibilities that could account for this tension, one particularly interesting property involves the existence of ***quasi-bound-state effects*** in the $t\bar{t}$ system, commonly referred to as "toponia". The formation of a colour-singlet quasi-bound $t\bar{t}$ state was first predicted in 1987 [18–20], nearly a decade before the discovery of the top quark at the Tevatron collider in 1995 [8, 9]. The threshold region has become increasingly important in the light of recent ATLAS [21] and CMS [22] results where an excess compatible with quasi-bound-states formation at threshold has been observed by both experiments.

Unlike the well-known bound states of charm and bottom quarks (charmonium and bottomonium), the top quark decays before it has time to hadronise due to its short lifetime. As a result, the finite decay width of the top quark prevents the formation of narrow, individually resolvable bound states. Instead, the contributions from multiple states near the threshold merge into a single broad enhancement in the $t\bar{t}$ invariant-mass spectrum. Since the top quarks decay before a fully formed bound state can develop, the term "quasi-bound state" is used, rather than simply "bound state". In this kinematic region, the top and antitop quarks have small relative velocities, and the dynamics of the $t\bar{t}$ system can therefore be described using the framework of non-relativistic QCD (NRQCD).

In the NRQCD regime, the interaction between the top and antitop quarks can be described by a Coulomb-like potential. This potential depends on the colour configuration of the $t\bar{t}$ system and can be either attractive or repulsive. In particular, colour-singlet configurations experience an attractive potential, while colour-octet configurations are repulsive. As a result, quasi-bound-state effects are present only in the colour-singlet channel, whereas they are suppressed in colour-octet configurations.

At the LHC, $t\bar{t}$ pairs can be produced in different colour and spin configurations, depending on the initial state. As discussed earlier, the dominant production mechanism near threshold is gg fusion. In this production mode, the $t\bar{t}$ pairs can be produced in colour-singlet or colour-octet states. In contrast, the subdominant $q\bar{q}$-initiated production proceeds predominantly through colour-octet configurations, and therefore does not contribute significantly to quasi-bound-state formation. Near threshold, the Coulombic interaction becomes significant compared to the kinetic energy of the top quarks. The characteristic interaction scale is of the order of $(\alpha_s m_t)^{-1}$, where α_s is the strong coupling constant and m_t is the top-quark mass. In the attractive colour-singlet channel, multiple-gluon exchange can therefore give rise to quasi-bound-state effects, leading to enhancements in the invariant-mass

spectrum just below the $t\bar{t}$ production threshold. This enhancement primarily originates from the attractive QCD potential, which lowers the energy of the system by the binding energy of the top quarks. An illustration of such a state is shown in Figure 3.10. At larger invariant masses, the top quarks become highly relativistic, and their kinetic energy dominates over the QCD potential, such that no appreciable binding occurs.

In the framework of NRQCD, the spin and orbital structure of the produced $t\bar{t}$ system is conveniently described using the LS (Russell–Saunders) coupling scheme [89], in which states are labelled by the term symbol $^{2S+1}L_J$. Here, S denotes the total spin of the $t\bar{t}$ pair, L their relative orbital angular momentum, and J the total angular momentum, satisfying $\vec{J} = \vec{L} + \vec{S}$. This notation provides a convenient way to classify the different configurations contributing near threshold, each of which exhibits distinct threshold behaviour and varying sensitivity to quasi-bound-state effects. Since the relative velocity of the top quarks is small near the production threshold, configurations with low oribital angular momentum dominate. The $t\bar{t}$ system is therefore predominantly produced in an S-wave ($L = 0$) configuration, while higher waves such as P-wave ($L = 1$) contributions are suppressed.

For gg fusion production, the dominant colour-singlet contribution near threshold corresponds to the $^1S_0$ configuration, characterised by a spin-singlet ($S = 0$) state and total angular momentum $J = 0$. This spin-singlet component corresponds to the maximally entangled spin configuration discussed earlier in the chapter. Depending on the degree of gluon exchange, multiple $^1S_0$ configurations with masses around the $t\bar{t}$ production threshold can contribute to the observed invariant-mass spectrum. Although $q\bar{q}$-initiated production can also contribute to spin-triplet configurations such as $^3S_1$, these contributions are expected to be subdominant near threshold, where gg fusion dominates. Therefore, near the $t\bar{t}$ production threshold at the LHC, the dominant NRQCD contribution to the $t\bar{t}$ cross-section originates from colour-singlet $t\bar{t}$ production via gg fusion, for which the attractive QCD potential can give rise to the quasi-bound states.

Figure 3.11 shows a comparison of the invariant mass distributions for $gg \to t\bar{t}$ colour-singlet (red) and colour-octet (blue) states in the $J = 0$ channel. It can be seen that the enhancement due to bound-state effects appears only in the colour-singlet channel, as expected. The dashed lines represent the cross-section summed over all possible J values. Further details on how these contributions are estimated can be found in Ref. [90].

A key challenge in measuring such quasi-bound-state effects at the LHC lies in the accurate modelling of the low-$m_{t\bar{t}}$ kinematic regime. Various theoretical efforts have been made to incorporate NRQCD effects [90, 91], including the resummation of soft and collinear gluon emissions near threshold [92, 93]. Such effects have so far not been included in the standard pQCD modelling of $t\bar{t}$ production. In addition to the missing NRQCD effects, higher-order QCD and electroweak corrections, as well as off-shell[5] effects, can also affect the description of $t\bar{t}$ production in the threshold regime.

In its entanglement analysis, CMS included quasi-bound-state effects by utilising a simplified model in which these effects are modelled as a pseudo-scalar s-channel resonance [94]. This choice is motivated by the fact that the dominant S-wave colour-singlet contribution to $gg \to t\bar{t}$ production arises mainly from the lightest pseudo-scalar $^1S_0^{[1]}$ configuration. Other J states, as well as the colour-octet contributions, are neglected in this model. Recently, progress has been made in the modelling to incorporate NRQCD effects to the $t\bar{t}$ threshold, leading to an improved model which includes all

[5] An on-shell particle satisfies the mass-shell condition $p^2 = m^2$, corresponding to a real, resonant particle with its physical mass. In contrast, a virtual (off-shell) particle does not satisfy this relation and appears as an internal propagator in a Feynman diagram, where its invariant mass can differ from the physical mass.

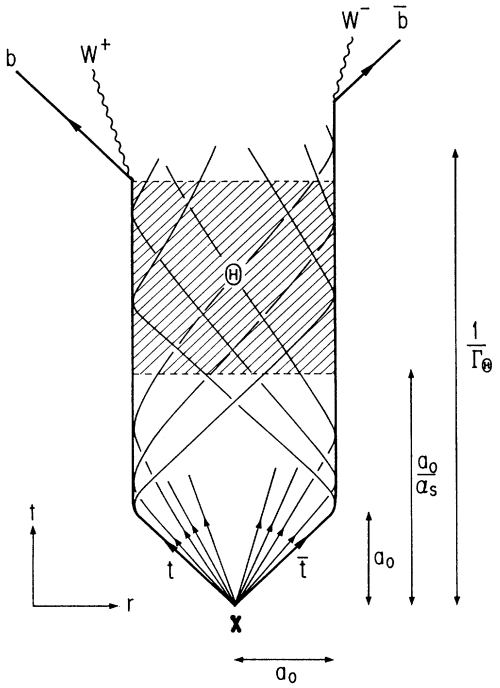


Figure 3.10: Illustration of the time evolution of a non-relativistic $t\bar{t}$ system leading to the formation of a quasi-bound-state via multiple-gluon exchange. The initial plane-wave components describing the $t\bar{t}$ system propagate outwards from the origin x=0. These waves evolve until they reach the effective potential barrier at a characteristic distance given by the Bohr radius, $a_0 = (\alpha_s m_t)^{-1}$. At times $t \sim a_0/\alpha_s$, the typical waves related to the bound states reach the barrier, and the formation of resonance-like quasi-bound configurations becomes relevant. The system remains in this state until the t or $\bar{t}$ decays at the timescale Γ_Θ^{-1}, the lifetime of this bound state [20].

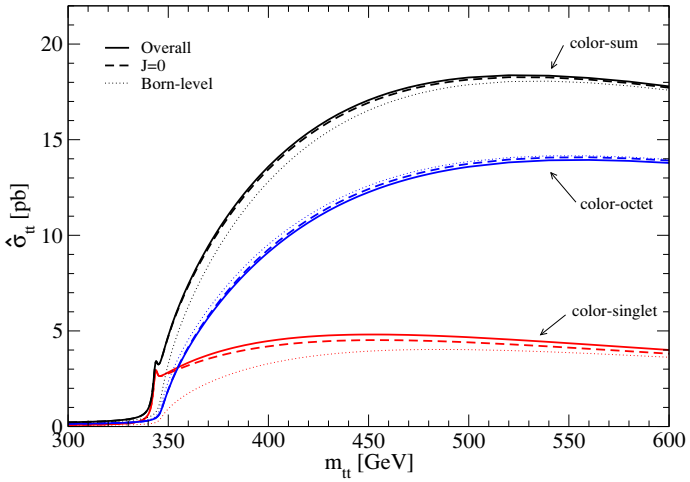


Figure 3.11: Distributions of the $t\bar{t}$ invariant mass for $gg \to t\bar{t}$ in the color-singlet (red), and color-octet (blue) states, and the sum of them (black). The solid lines represent the $J = 0$ contribution and dashed lines are calculated as the sum of contributions from all J values [90].

S-wave contributions spanning the full spectrum of ${}^1S_0$ mass states around the $t\bar{t}$ production threshold. This includes effects from both bound-state contributions and non-resonant scattering states [95]. However, effects from subdominant $q\bar{q}$-initiated contributions to ${}^3S_1$ states, higher waves (such as P-wave) and from colour-octet states are not included. This improved model based on NRQCD $t\bar{t}$ production at threshold provides a more complete description of $t\bar{t}$ quasi-bound-state formation than the simplified model.

The latest CMS measurement at threshold [22] uses the simplified model, while the ATLAS measurement [21] used the improved model, with results also reported using the simplified model for comparison with CMS. The present work uses the improved model to incorporate quasi-bound-state effects. This model has only recently become available for the lepton+jets channel [95], making it the first entanglement measurement in this channel to incorporate this modelling. Results obtained using the simplified pseudo-scalar resonance model are also presented for comparison. Details of the corresponding simulated samples are provided in Chapter 5.

CHAPTER 4

The ATLAS experiment at the LHC

"My first visit to the Large Hadron Collider was in April 2008, before it started up, and CERN had some open days. They were slightly shocked by the end of it because I think they got something like 50,000 visitors."

— Peter Higgs

The discovery of the top quark at the Tevatron collider at Fermilab in 1995 [8, 9] marked a major milestone in particle physics and completed the current SM quark sector. Since the Tevatron accelerated protons and antiprotons to energies of the order of 1 TeV, it enabled the production of large numbers of top-quark events. In general, detailed studies of top-quark properties benefit from experimental conditions capable of producing high-energy collisions that provide sufficient top-quark production rates. One such facility is the ***Large Hadron Collider (LHC)*** [96] at CERN, currently the most powerful particle accelerator and collider in the world. Built on the legacy of previous machines such as HERA and the Tevatron, the LHC provides high energies and collision rates to study the top-quark sector with unprecedented precision. These high energies result in large top-quark production cross-sections (approximately 800 pb[1] for $t\bar{t}$ production), effectively making the LHC a top-quark factory. This enables detailed studies of top-quark properties across a wide kinematic range and even provides access to rare processes such as four-top-quark production.

This chapter presents a general overview of the LHC, including its structure, operation, and experimental detectors, with particular emphasis on the ***ATLAS experiment***, one of the two general-purpose detectors at the collider.

The chapter opens with an overview of the LHC complex and its operation. Section 4.1.1 describes the key collider parameters relevant to the analyses presented in this thesis. Section 4.2 introduces the ATLAS experiment and its subdetectors, describing the data acquisition process, the nomenclature used within the ATLAS collaboration for different data-taking periods, and the procedures employed to assess data quality. ATLAS continues to operate currently, with a long shutdown planned to start in the summer of 2026 to implement the major upgrades required for the high-luminosity LHC era. These upgrade plans are discussed in Section 4.3.

[1] The prefix "pico" (p) denotes a factor of 10^{-12}. Thus, 1 pb = 10^{-12} b.

4.1 The Large Hadron Collider

The LHC is located at CERN, near Geneva, and lies in a circular underground tunnel with a circumference of 27 km, at depths ranging from 50 to 100 m below the surface. Primarily designed to conduct proton–proton (pp) collisions, the LHC is also capable of colliding heavy ions or atomic nuclei, enabling interactions such as lead–lead, proton–lead, oxygen–oxygen, and proton–oxygen collisions. For the purpose of this thesis, only pp collision data are considered. At the LHC, protons are accelerated and guided in the form of tightly focused beams. The machine consists of two parallel beam pipes, each containing a beam circulating in opposite direction. The major systems ensuring its successful operation include the vacuum system, superconducting magnets, cryogenics, and radio-frequency (RF) cavities that accelerate and stabilise the beams. The overall structure of the accelerator complex is shown in Figure 4.1.

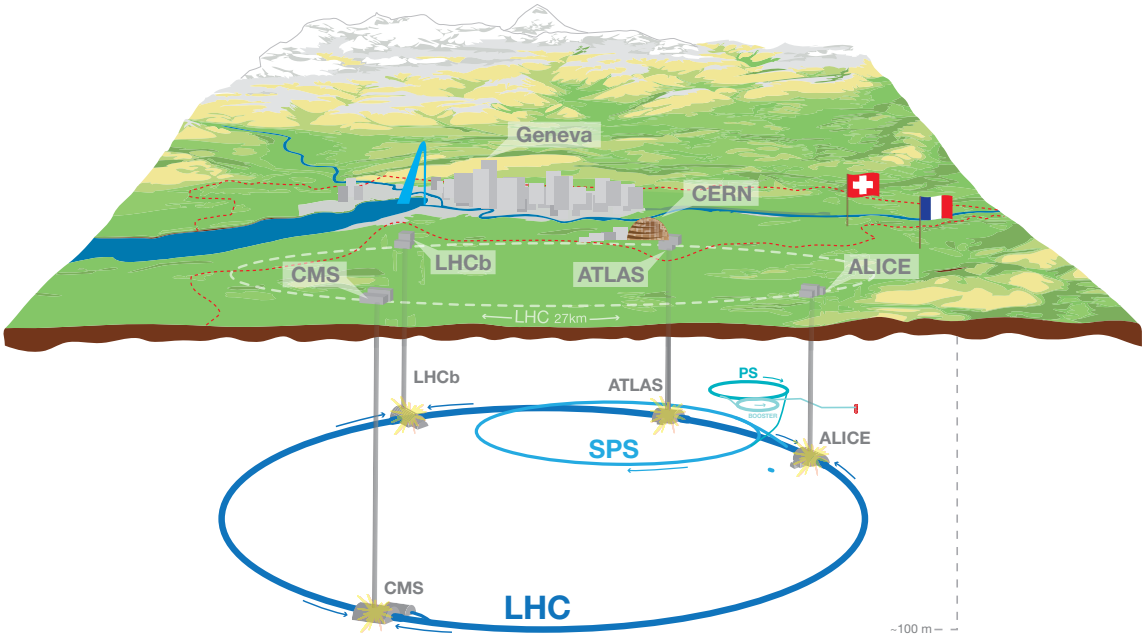


Figure 4.1: A schematic illustration of the LHC complex. The four large-scale detectors are shown at the four main IPs, together with the pre-accelerator chain (BOOSTER, PS, and SPS), which boosts the protons before injecting them into the main LHC ring [97].

The proton beams are not continuous but are grouped together in distinct bunches, with each ***bunch*** containing $\approx 10^{11}$ protons. The bunch spacing is typically 25 ns, set by the RF frequency used for acceleration. Collisions occur when two bunches from opposing beams pass through each other, an event referred to as a ***bunch crossing***. The beams are guided around the ring by 8 T superconducting ***dipole magnets***, while ***quadropole magnets*** focus and defocus them to maintain beam stability. These magnets are superconducting electromagnets operating at 1.9 K in baths of superfluid helium. At such low temperatures, they conduct currents with negligible resistance, thus saving energy from being lost to electrical resistance. The LHC ***cryogenic system*** maintains these temperatures through a

network of refrigerators and cold compressors distributed around the ring. A dedicated cryogenic ring supplies the necessary cooling power to the magnets. The ***vacuum system*** of the LHC is equally crucial, providing an ultra-clean environment to prevent beam particles from interacting with residual gas molecules, which could degrade beam quality or damage the accelerator.

The protons for collisions are extracted from a small bottle of hydrogen gas, where the electrons are stripped off using strong electric fields. They are then accelerated through a series of pre-accelerators before injection into the LHC, as illustrated in Figure 4.1. They are first accelerated by the linear accelerator (LINAC 4[2]) to 50 MeV, then boosted to 25 GeV by the Proton Synchrotron Booster (BOOSTER) and Proton Synchrotron (PS). The Super Proton Synchrotron (SPS) further accelerates them to 450 GeV, at which point they are injected into the LHC. Inside the LHC, the beam is split into two counter-rotating beams that are then accelerated to their peak energy of 6.5 TeV each[3]. These beams collide at four main ***interaction points (IPs)***, each hosting one of the large-scale LHC experiments: ATLAS[4] [98], CMS[5] [99], ALICE[6] [100], and LHCb[7] [101]. The ALICE experiment focuses on the study of the quark-gluon plasma through heavy-ion collisions, while LHCb is optimised for studies involving b-hadrons. ATLAS and CMS are the two general-purpose detectors designed for broad studies of the SM and searches for new physics phenomena. In addition to these large-scale detectors, the LHC also hosts several smaller experiments dedicated to specialised measurements and studies. This thesis focuses on the ATLAS experiment, whose structure and functionality are described in detail in the following section. Before that, it is useful to introduce several key collider parameters commonly used throughout this work.

4.1.1 Essential collider parameters

Centre-of-mass energy

The ***centre-of-mass energy***, mentioned previously in Section 3.1.1, is a measure of the total available energy in the system of colliding particles. It effectively defines the upper energy threshold for the production of new particles. Higher centre-of-mass energies allow access to more massive states and increase the probability of producing rare processes. It is conventionally denoted by $\sqrt{s}$, where s is one of the Mandelstam variables, representing the square of the total energy in the centre-of-mass frame. Hence, $\sqrt{s}$ corresponds to the total energy available for particle production in the collision. At the LHC, pp collisions are organised into distinct operational periods, with each period known as a ***Run***. The first operation period, Run 1, lasted from 2010 to 2012, during which proton beams were collided at $\sqrt{s} = 7$ TeV and later increased to 8 TeV in 2012. After a two-year long technical shutdown, the LHC resumed operations from 2015 to 2018 for Run-2, operating at $\sqrt{s} = 13$ TeV. This thesis is based on Run-2 data. The most recent period, Run-3, began in 2022 and is expected to continue until 2026, with collisions occurring at an increased centre-of-mass energy of $\sqrt{s} = 13.6$ TeV. Following

[2] LINAC 4 is the successor of LINAC 2, which was used during the 2015–2018 data-taking period, referred to as Run-2. The ongoing data-taking period is Run-3 (Section 4.1.1).

[3] Note that this value corresponds to the peak beam energy during the Run-2 data-taking period. The beam energy can vary between different runs.

[4] **A** **T**oroidal **L**HC **A**pparatu**S**

[5] **C**ompact **M**uon **S**olenoid

[6] **A** **L**arge **I**on **C**ollider **E**xperiment

[7] **LHC**-**b**eauty

Run-3, the LHC is expected to enter a long shutdown period in preparation for its upgrade to the high-luminosity LHC (Section 4.3).

Luminosity

While the centre-of-mass energy in pp collisions determines the range of accessible physics processes, it does not directly control the frequency at which a particular process occurs. Another key parameter that quantifies the performance of the collider in this regard is the *instantaneous luminosity*, defined as:

$$\mathcal{L}_i = \frac{f n_1 n_2}{4\pi \sigma_x \sigma_y}, \tag{4.1}$$

where f is the bunch-crossing frequency in the accelerator ring, n_1 and n_2 are the number of particles per colliding bunch, σ_x and σ_y are the horizontal and vertical beam sizes at the IP. The event rate for a given process is proportional to the product of the instantaneous luminosity $\mathcal{L}_i$ and the process cross-section σ. The expected number of events over a data-taking period T, assuming a constant cross-section, is therefore given by:

$$N = \sigma \int_0^T \mathcal{L}_i(t)\, dt = \sigma \mathcal{L},$$

where $\mathcal{L}$ is referred to as the ***integrated luminosity***. It represents the total delivered luminosity over the period of interest and determines the total number of events available for analysis. In the particle physics community, it is commonly expressed in units of inverse barn (b^{-1}), where $1\,\mathrm{b}^{-1} = 10^{28}\,\mathrm{m}^{-2}$. While the cross-section is a property of the underlying physics process, luminosity characterises the machine performance. For searches involving rare processes, maximising $\mathcal{L}$ over long data-taking periods is essential to increase the expected number of observable events.

The luminosity measurement in ATLAS relies on several subdetectors and algorithms, each providing complementary capabilities and independent systematic uncertainties [102]. The primary detectors used during Run-2 were the LUCID-2 Cherenkov Detector [103] and the Beam Conditions Monitor (BCM) [104]. LUCID-2, positioned symmetrically at a distance of 17 m from the ATLAS IP, provides the main bunch-by-bunch luminosity measurement, complemented by measurements from the BCM, which consists of four $8 \times 8\,\mathrm{mm}^2$ diamond sensors located at 1.84 m on either side of the IP. In addition to its role in luminosity determination, the BCM also monitors beam conditions to protect the detector from potential damage due to beam losses. These detectors along with various other offline measurements detailed in Ref. [102] provide the relative luminosity measurements, i.e. variations in luminosity over time. These measurements are then related to the absolute luminosity through a calibration procedure known as the van der Meer (vdM) scans [105].

The vdM method scans the overlap of the two beams in the transverse plane ($x - y$ plane) and measures the interaction rates as a function of their separation. The horizontal and vertical planes are scanned independently, allowing a complete determination of the effective transverse beam sizes $\sigma_x \sigma_y$ that appear in Equation 4.1.

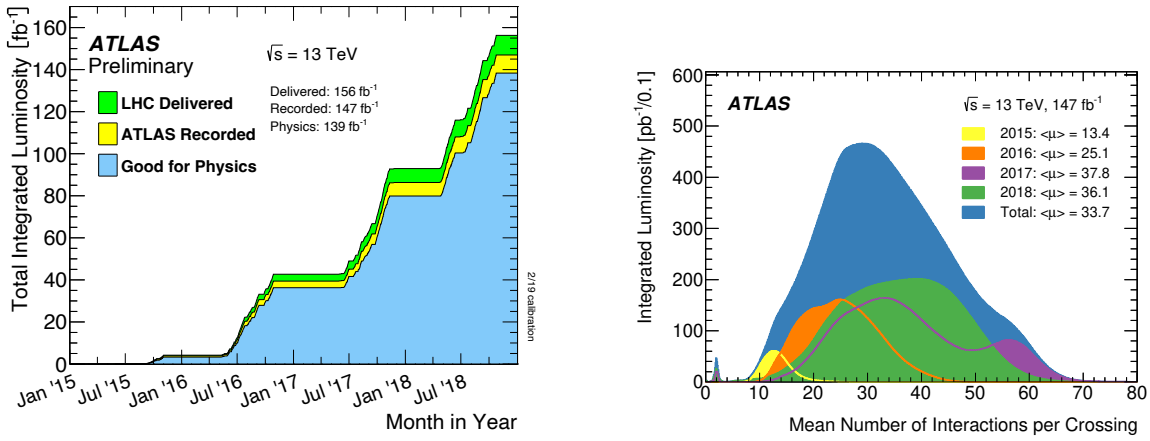
The integrated luminosity delivered to the ATLAS experiment during Run-2 is shown in Figure 4.2(a). During Run-2, the LHC delivered $\approx 156\,\mathrm{fb}^{-1}$ of integrated luminosity[8], of which $\approx 147\,\mathrm{fb}^{-1}$ was

[8] The prefix "femto" (f) denotes a factor of 10^{-15}. Thus, $1\,\mathrm{fb}^{-1} = 10^{15}\,\mathrm{b}^{-1}$.

recorded by the ATLAS detector. The slight reduction arises from detector efficiencies and data-quality requirements (Section 4.2.6) that determine which events are suitable for physics analyses. The total integrated luminosity used for the Run-2 analyses in this thesis corresponds to $140.1 \pm 1.2\,\text{fb}^{-1}$ [102]. As of May 2026, the ongoing Run-3 of the LHC has delivered $\approx 351\,\text{fb}^{-1}$ of integrated luminosity, representing one of the largest datasets collected in collider physics, of which $\approx 332\,\text{fb}^{-1}$ has been recorded by the ATLAS detector [106].

Pile-up

The high instantaneous luminosity comes at the cost of a phenomenon known as ***pile-up***, which refers to the presence of multiple pp interactions either within a single or adjacent bunch crossing. There are two main types of pile-up. *In-time pile-up* refers to all pp interactions occurring in the same bunch crossing as the primary hard interaction of interest. *Out-of-time pile-up* comes from detector signals produced by interactions in bunch crossings immediately before or after the one of interest, due to the finite signal-formation time and the finite duration of signals in the detector. The amount of pile-up in collision data is typically quantified by the average number of interactions per bunch crossing, $\langle\mu\rangle$. The distribution of μ is shown in Figure 4.2(b) for the IP of the ATLAS experiment during Run-2. For the full Run-2 pp collision dataset, an average of 33.7 interactions per bunch crossing was recorded. For Run-3, which is expected to conclude in 2026, the corresponding value had reached approximately 56 as of 2025, representing an increase of about 66% relative to Run-2. Mitigating pile-up effects is one of the major challenges for physics analyses at the LHC, and uncertainties related to the modelling of the pile-up profile in simulations are typically included in the measurement (Section 7.9).



(a) (b)

Figure 4.2: (a) Integrated luminosity as a function of time delivered to ATLAS (green), recorded by ATLAS (yellow), and passing data-quality criteria (blue) [106] (b) Distribution of the mean number of interactions per bunch crossing, $\langle\mu\rangle$, at the ATLAS experiment for pp collisions during Run-2. The quantity $\langle\mu\rangle$ is calculated from the instantaneous luminosity assuming an inelastic pp collision cross-section of 80 mb at $\sqrt{s} = 13\,\text{TeV}$ [107].

4.2 The ATLAS experiment

The ATLAS detector [98, 108], shown in Figure 4.3, is a large, cylindrical, multipurpose particle detector surrounding the LHC beamline at one of the IPs. It is designed primarily to study pp collisions, but is also used for heavy-ion physics and for studies using proton–oxygen collisions. The detector is approximately 46 m long, 25 m in diameter, and weighs about 7000 tonnes. Its geometry is optimised to provide nearly full solid-angle coverage around the IP, ensuring that most particles produced in hard-scattering interactions traverse active detector components. The central cylindrical region is referred to as the ***barrel***, while the two sections covering the forward regions along the beamline are known as the ***end-caps***.

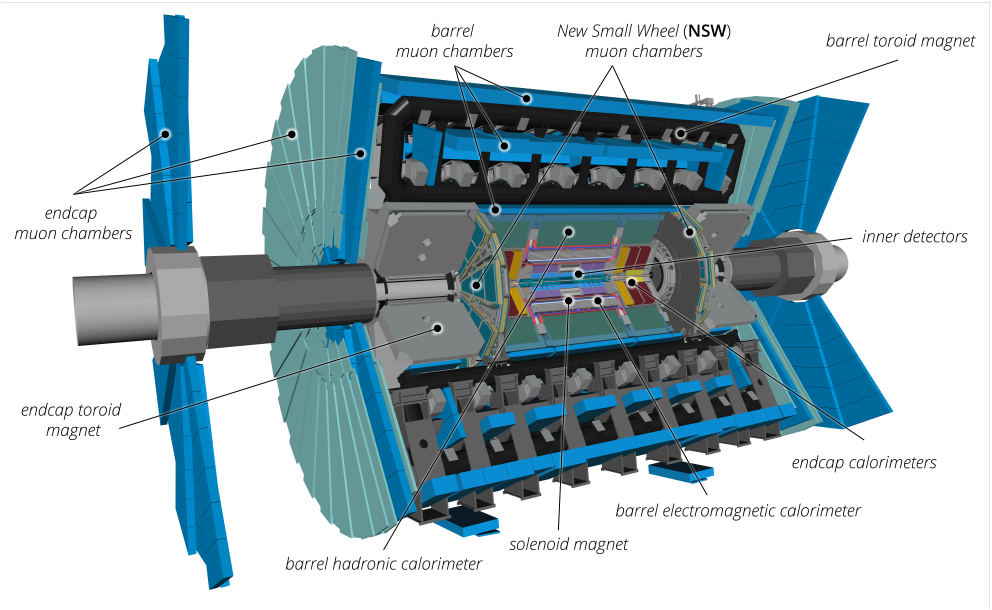


Figure 4.3: Cut-away view of the ATLAS detector [108], shown in its Run-3 configuration, which differs from the Run-2 setup primarily through upgrades to the Small Wheels (Section 4.2.4).

A cartesian coordinate system is used for measurements in the detector, with the origin located at the pp collision point. The beam pipe lies along the z-axis while the x-y plane is transverse to the beam direction. Cylindrical coordinates are also commonly used: the ***azimuthal angle***, ϕ, is measured around the beam axis, and the ***polar angle***, θ, is defined with respect to the beam direction. From θ one can define the ***pseudorapidity***, η, a commonly used variable describing the particle direction relative to the beam pipe:

$$\eta = -\ln\left(\tan\left(\frac{\theta}{2}\right)\right).$$

The pseudorapidity is zero in the transverse plane ($\theta = \pi/2$) and increases rapidly in magnitude near the beam axis with $|\eta| \to \infty$ for $\theta \to 0$ or π. In the ultra-relativistic limit, where the particle mass is negligible compared to its momentum ($m \ll p$) and the total energy satisfies $E \approx p$, the

pseudorapidity equals the ***rapidity*** y, defined as :

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right).$$

Here, p_z is the z-component of the particle momentum. Using pseudorapidity instead of the angle θ has the advantage that differences in pseudorapidity, $\Delta\eta = \eta_1 - \eta_2$, are invariant under Lorentz boosts along the beam axis. This is useful because at the LHC, the centre-of-mass frame of a process does not necessarily coincide with the laboratory frame. The momentum fractions of the initial-state partons, and thus their longitudinal momenta, vary event by event according to the PDFs. In contrast, the transverse momentum components of the partons are approximately zero. Momentum conservation requires the vector sum of the transverse momenta of all final-state particles to be zero. By measuring the transverse momenta of all particles that interact in the detector and taking their vector sum, one can determine the ***missing transverse momentum*** or ***missing transverse energy***, $E_{\mathrm{T}}^{\mathrm{miss}}$ (Section 6.5).

Based on this, it is common practice to characterise the four-momentum $\vec{p}$ of a particle not by its cartesian components (p_x, p_y, p_z, E) but by (p_T, η, ϕ, E). Here $p_{\mathrm{T}} = \sqrt{p_x^2 + p_y^2}$, and equivalently $p_{\mathrm{T}} = |\vec{p}| \sin\theta$. Since differences in azimuthal angle $\Delta\phi$ and pseudorapidity $\Delta\eta$ are invariant under Lorentz boosts along the direction of the beam axis, a distance variable can be defined in the $\eta - \phi$ plane:

$$\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$$

In physics analyses, ΔR is used as the general distance dimension, often referred to as the angular separation. It is also invariant under Lorentz boosts along the z-axis. An illustration of the coordinate system is shown in Figure 4.4.

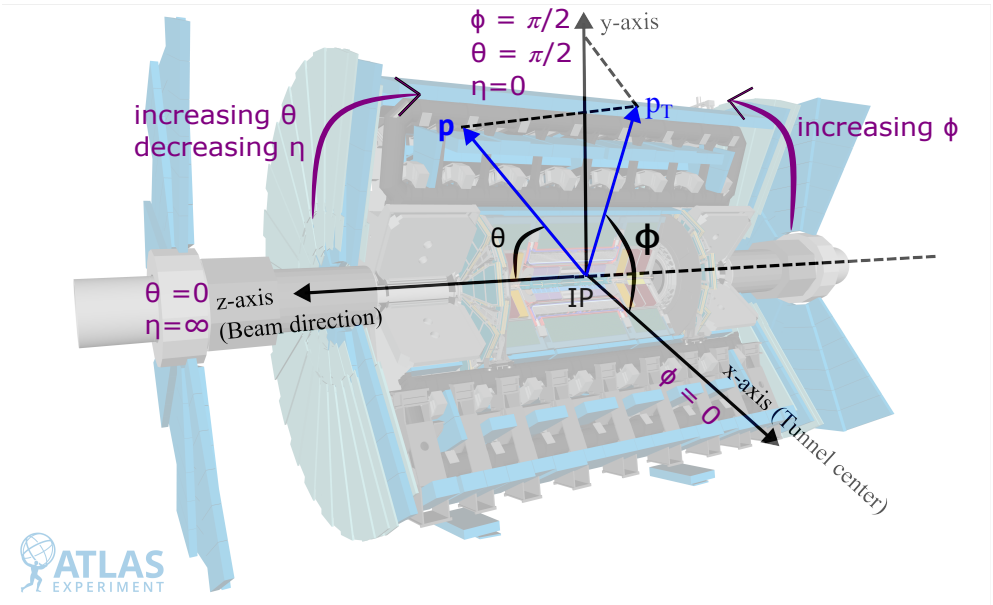


Figure 4.4: An illustration of the ATLAS coordinate system. Image taken from Ref. [109] and modified.

The main components of the ATLAS detector, moving outward from the IP are the Inner Detector,

used to measure trajectories of charged particles (Section 4.2.1), the calorimeters, which measure the energies of most charged and neutral particles (Section 4.2.2), and the muon spectrometer, specialised for detecting muons (Section 4.2.4). A magnet system (Section 4.2.3) provides the magnetic fields required to bend the trajectories of charged particles to determine their momentum and charge. Finally, the trigger and data acquisition system (Section 4.2.5) manages the enormous data rates produced by the detector and records only the most interesting events for physics analysis (Section 4.2.6).

4.2.1 Inner Detector

The ***Inner Detector (ID)*** is the innermost component of the ATLAS detector and is therefore closest to the beam pipe. It primarily measures the trajectories of charged particles by recording their intersections with the active detector layers. The resulting reconstructed trajectory is referred to as a ***track***. The ID is 6.2 m long, 2.1 m in diameter, and operates inside a 2 T axial magnetic field, which bends charged-particle tracks and allows their momenta to be measured in the transverse plane (Section 4.2.3). Being the innermost subsystem, it is highly sensitive to the large particle flux near the beam line. To cope with the high track density expected at the LHC, the ID employs very high-granularity detectors, providing full ϕ coverage and acceptance of up to $|\eta| < 2.5$. The granularity is highest near the interaction, enabling precise reconstruction of the ***primary vertex (PV)***, defined as the vertex with the largest sum of the scalar transverse momenta of its associated tracks. Accurate PV identification is essential, since it is used to select the interaction corresponding most likely to the hard scatter and hence the physics process of interest. The fine spatial resolution of the ID also aids in rejecting tracks from pile-up vertices, which are typically displaced along the z-axis. The ID is also crucial for reconstructing secondary vertices from decays of short-lived particles such as B- and D-hadrons. ***Secondary-vertex*** information is an important discriminant in flavour-tagging algorithms (Section 6.4.6).

The requirements of the tracking system vary with the distance from the IP. Close to the IP, high-granularity, radiation-tolerant detectors are required for efficient vertexing. At larger distances, the particle flux is lower and the spatial resolution requirements are reduced. Therefore, different detector technologies are used to meet these needs in a cost-effective manner, resulting in three main ID subsystems: the ***Silicon Pixel Detector***, the ***Semiconductor Tracker (SCT)***, and the ***Transition-Radiation Tracker (TRT)***, as shown in Figure 4.5. The pixel detector and SCT use silicon sensors, while the TRT uses gas-filled straw tubes.

Pixel Detector The Pixel Detector is the closest subsystem to the IP and is a silicon-based semiconductor tracker consisting of four layers in the barrel and three disks in each end-cap. Its sensors are segmented into a grid of rectangular elements, referred to as ***pixels*** [110], with dimensions of 50 μm in the transverse direction and 250 – 400 μm along the beamline. When a charged particle passes through the detector, it ionises the active silicon material, depositing charge that is collected by the nearby pixels. Each charge deposit is recorded as a hit. Within the acceptance of the tracking system ($|\eta| < 2.5$), particles typically traverse four pixel layers. The most recent addition (installed for Run-2) is the innermost layer, known as the insertable B-layer (IBL) [111] (Figure 4.5). The IBL was added to improve tracking and vertexing performance and to mitigate radiation damage effects on the original innermost pixel layer.

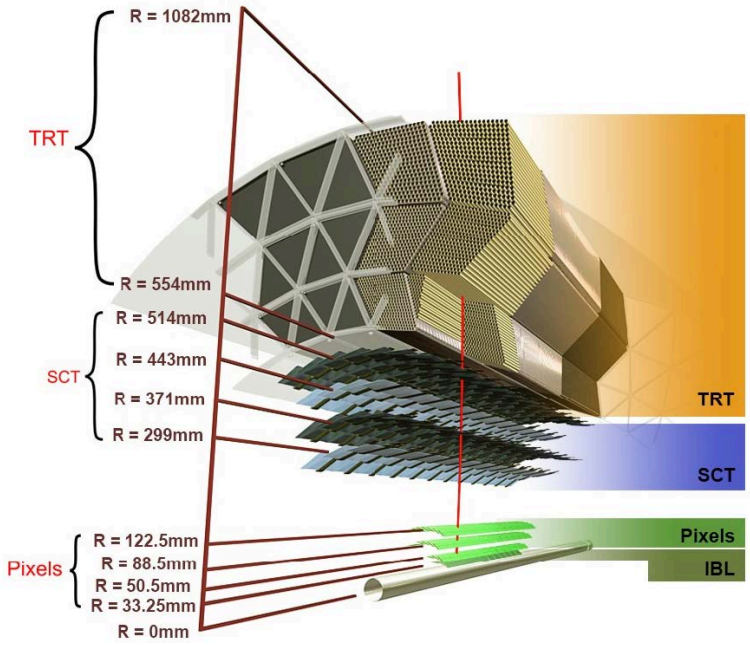


Figure 4.5: Transverse section of the ATLAS ID. Going outward from the beamline, the ID subsystems are: the Silicon Pixel Detector, the Semiconductor Tracker (SCT), and the Transition-Radiation Tracker (TRT) [108].

Semiconductor Tracker The SCT is similar in concept to the pixel detector but uses long, thin silicon microstrips instead of pixels, with a pitch of about 80 μm [98]. In the barrel region, it spans radii from 299 to 514 mm and consists of four coaxial cylindrical layers, while each end-cap contains nine disk layers. Together, these provide at least four position measurements for tracks within the tracking acceptance. The SCT achieves a spatial resolution of approximately 17 μm in the transverse direction and 580 μm in the z-direction. A single microstrip layer provides position measurement only in the coordinate perpendicular to the strip orientation. To obtain two-dimensional hit positions, each SCT module is built from two back-to-back strip layers tilted by a small stereo angle.

Transition-Radiation Tracker The TRT is the outermost subsystem of the ID, covering radii from 554 to 1082 mm in the barrel region, with a reduced acceptance of $|\eta| < 2$ compared to the silicon-based trackers. It consists of roughly 300000 straw tubes (50000 in the barrel and 250000 in the end-cap), each 4 mm in diameter and filled with a xenon-based gas mixture, arranged in up to 73 layers in the barrel and 160 planes in the end-caps [112]. The straws are interleaved with polypropylene fibres (barrel) or foils (end-caps). When an ultra-relativistic charged particle crosses the polypropylene layer/fibre, it emits X-ray photons which are absorbed by the xenon gas in the straw, producing additional ionisation electrons. This phenomenon is known as ***transition radiation***. Each straw contains a 0.03 mm diameter gold-plated tungsten readout wire that collects these signals. Charged particles typically leave many hits in the TRT, improving the momentum measurements, and its large radius enhances the determination of track curvature. The detector, however, does not provide position information along the axis of the straw. Since electrons tend to emit more transition radiation than

pions[9], the TRT can use this effect to discriminate between the two particle types.

4.2.2 Calorimeters

The calorimeter system, illustrated in Figure 4.6, surrounds the ID and the solenoid (part of the magnet system discussed in Section 4.2.3), providing nearly full solid-angle coverage of up to $|\eta| = 4.9$. It is designed to perform destructive measurements of the energy of most charged and neutral particles, with the exception of muons and neutrinos, which traverse it with minimal interaction. The large angular coverage enables the reconstruction of the missing transverse momentum (Section 6.5).

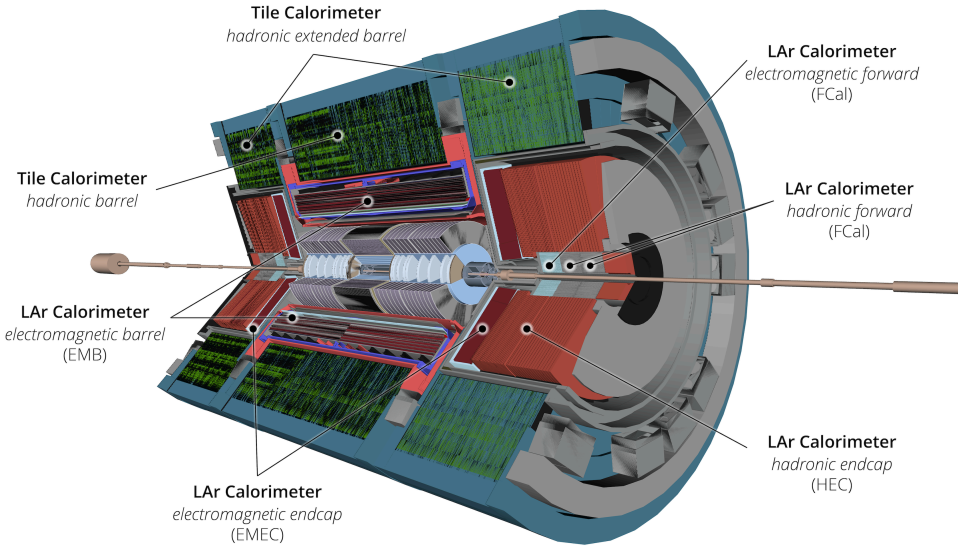


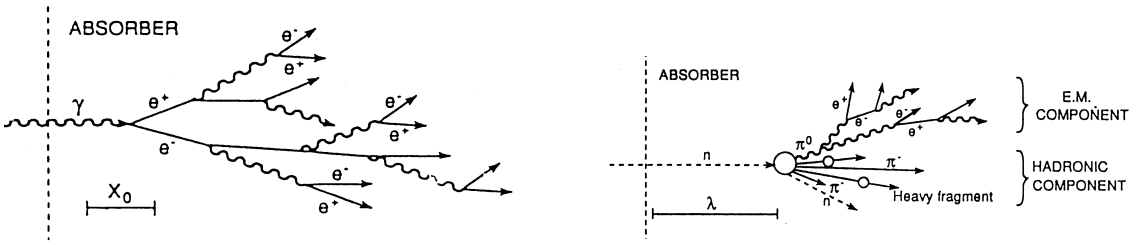
Figure 4.6: Illustration of the ATLAS calorimeter system. The electromagnetic calorimeter surrounds the ID and is overlaid by the hadronic calorimeter. The main components are labelled in the figure [108].

The calorimeter system consists of two main components. The ***electromagnetic calorimeter (ECal)***, located just outside the solenoid, measures the energies of electrons and photons. Surrounding it is the ***hadronic calorimeter (HCal)***, which measures the energies of charged and neutral hadrons. When a photon travels through the calorimeter, it produces an electron-positron pair. These electrons and positrons, as well as electrons produced in the collision, undergo a bremsstrahlung process and radiate photons. The radiated photons can subsequently convert into additional electron-positron pairs generating a self-sustaining cycle of bremsstrahlung and pair production. This process produces a cascade of secondary particles known as an ***electromagnetic shower***, as shown in Figure 4.7(a). The characteristic length scale of this process is the radiation length X_0, defined as the average distance over which an electron loses a factor of $1/e$ of its energy due to bremsstrahlung. The ECal dimensions are chosen to fully contain electromagnetic showers.

Hadrons interacting with the calorimeter material also produce cascades of secondary particles known as ***hadronic showers***. These arise from inelastic strong interactions between the hadrons and

[9] Pions (π mesons) are the lightest mesons, composed of quark–antiquark pairs, and appear in neutral (π^0) and charged ($\pi^{\pm}$) states.

atomic nuclei. These cascades contain both hadronic and electromagnetic components, the latter resulting mainly from decays such as neutral-pion decays, $\pi^0 \rightarrow \gamma\gamma$ (Figure 4.7(b)). The characteristic scale for this process is the interaction length λ, defined as the average distance a high-energy hadron travels in the calorimeter before undergoing a nuclear interaction. Since hadronic showers develop over greater distances than electromagnetic showers, the HCal has a radial thickness roughly four times that of the ECal. Some hadronic shower particles, including muons and neutrinos, may escape the calorimeter, and some electromagnetic shower components may leak into the HCal. Such an effect is called a *punch-through* and these are accounted for through specific calibration procedures for hadronic jet energy reconstruction (Section 6.4). Both the ECal and HCal employ a sampling calorimeter design, consisting of alternating layers of absorber material to induce shower development and active material that samples the deposited energy as the shower evolves.



(a) (b)

Figure 4.7: (a) Illustration of an electromagnetic shower initiated by a high-energy photon. (b) Schematic depiction of a hadronic shower showing its electromagnetic and hadronic components [113].

The electromagnetic calorimeter The ECal is divided into a barrel region ($|\eta| < 1.475$), end-caps ($1.375 < |\eta| < 3.2$), and a forward region ($3.1 < |\eta| < 4.9$) [114]. Its thickness exceeds 24 X_0 in the barrel and 26 X_0 in the end-caps, providing sufficient depth to contain electromagnetic showers. The barrel and end-cap calorimeters use liquid argon (LAr) as the active medium, interleaved with lead absorber plates bent into an accordion geometry (Figure 4.8), ensuring full azimuthal coverage. In the central region, $|\eta| < 2.5$, which is used for most physics analyses, the ECal is segmented into three longitudinal layers. The first layer is finely segmented to provide an accurate measurement of η, improving particle identification. In particular, it helps distingush a single photon γ from the collimated photon pairs produced in neutral-pion decays ($\pi^0 \rightarrow \gamma\gamma$), where the two photons tend to have a small angular separation in the laboratory frame. The middle and third layers have granularities of $\Delta\eta \times \Delta\phi = 0.025 \times 0.0245$ and 0.05×0.0245, respectively. An additional layer called the *presampler layer*, placed in front of the first layer, is used to correct for energy lost by particles before entering the calorimeter. The forward part of the calorimeter system ($3.1 < |\eta| < 4.9$), known as the ***forward calorimeter (FCal)***, includes an electromagnetic section that uses copper absorbers instead of lead and consists of a single longitudinal layer, optimised for operation in the high-radiation, high-occupancy forward region.

The hadronic calorimeter The HCal consists of a barrel tile calorimeter (TileCal) and the hadronic end-cap calorimeters (HEC). The total thickness is 11 λ at $\eta = 0$, providing good containment of hadronic showers and reducing punch-through into the muon system [114]. The TileCal covers

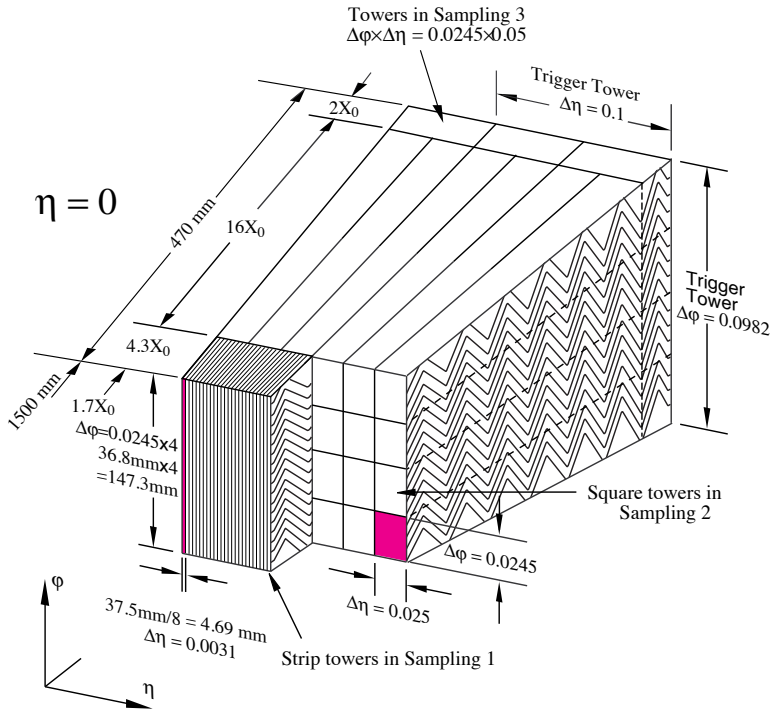


Figure 4.8: Sketch of the structure of an ATLAS electromagnetic LAr calorimeter barrel module [98].

$|\eta| < 1.7$ and uses iron as the absorber and plastic scintillating tiles as the active medium. The HEC covers the range $1.5 < |\eta| < 3.2$ and uses LAr as the active medium and copper plates as absorbers, ensuring robustness against the high radiation doses in the end-cap region. The hadronic section of the FCal, located directly behind the ECal section, uses LAr as the active medium and tungsten as the absorber. It is segmented into two layers. The FCal is designed to reconstruct energies of forward jets and, together with the HCal, contributes to a precise measurement of the missing transverse momentum.

4.2.3 Magnet system

The magnetic system of the detector [115] is illustrated in Figure 4.9 and consists of four large superconducting magnets: a ***central solenoid (CS)*** surrounded by three large air-core toroids, comprising one ***barrel toroid (BT)*** and two ***end-cap toroids (ECTs)***. The CS provides a magnetic field of about 2 T along the z-axis in the ID, bending charged-particle tracks in the transverse plane and enabling transverse momentum measurements in the ID.

The BT and ECTs provide the magnetic field for the muon spectrometer (Section 4.2.4), reaching peak field strengths of 3.9 T and 4.1 T, respectively. Each toroid consists of eight rectangular coils arranged symmetrically around the beam axis. The ECT coil system is rotated by 22.5° relative to the BT coil system to optimise the bending power in the transition region between the barrel and end-cap. The entire magnet system is approximately 26 m long and 22 m in diameter, and is maintained at cryogenic temperatures [114].

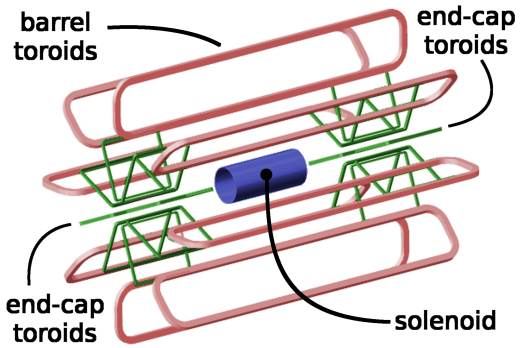


Figure 4.9: A schematic view of the ATLAS magnet system: the solenoid (blue), the barrel toroid (pink) and the two end-cap toroids (green) [116].

4.2.4 Muon Spectrometer

The ***Muon Spectrometer (MS)***[10] [108] is the outermost subsystem of the ATLAS detector. It is designed to measure the momenta of muons exiting the calorimeters with high precision over the pseudorapidity range $|\eta| < 2.7$. The momentum resolution can be further improved by combining MS information with that from the ID (Section 6.1). The magnetic bending is provided primarily by the large BT, supplemented by two ECTs covering $1.4 \leq |\eta| \leq 2.7$ and $1.0 \leq |\eta| \leq 1.4$. This toroidal configuration produces a magnetic field that is mostly orthogonal to muon trajectories.

Since the MS spans a very large volume, continuous tracking is not feasible. Instead, muon tracks are reconstructed from straight segments measured in three spatially separated stations [114]. The inner station lies just outside the calorimeters and upstream of most of the toroidal field. The middle station is situated inside or immediately downstream of the toroid, where the bending is the largest. The outer station is positioned at the largest radius, outside the magnetic field. Curved tracks are obtained by fitting through these three measurement points and are subsequently matched to tracks reconstructed in the ID. The MS is segmented into three cylindrical layers in the barrel and three disk-shaped wheels in each end-cap. These stations contain two classes of detector technologies: precision tracking chambers and trigger chambers, as shown in Figure 4.10. Precise measurements of the muon trajectory are provided mainly by the ***monitored drift tubes (MDTs)***, which cover most of the acceptance ($|\eta| < 2.7$). Each MDT consists of aluminium tubes with a central Tungsten Rhenium wire and filled with a mixture of 93% argon and 7% CO_2. In the very forward region ($1.0 < |\eta| < 2.7$), the reconstruction is done by ***cathode strip chambers (CSCs)*** arranged in four concentric disks around the beam axis. The end-cap region is equipped with three muon chamber wheels perpendicular to the beam pipe: one small wheel upstream of the end-cap toroid and two larger wheels downstream, each consisting of multiple tracking planes. The trigger chambers comprise the ***resistive plate chambers (RPCs)*** in the barrel ($|\eta| < 1.05$) and three layers of ***thin gap chambers (TGCs)*** in the end-caps ($1.05 < |\eta| < 2.4$). These chambers deliver fast timing and bunch-crossing identification, enabling the ATLAS trigger system (Section 4.2.5) to decide whether to record an event. The geometrical acceptance of the trigger chambers system is about 80% in the barrel and 99% in the end-caps. The lower acceptance in the barrel arises mainly from unavoidable uninstrumened gaps for services and

[10] Unless stated otherwise, all detector descriptions correspond to the Run-2 configuration. In particular, the inner end-cap wheels refer to the Run-2 small wheels and not the new small wheels (NSWs) installed for Run-3.

support structures.

For Run-3, the innermost end-cap wheels were replaced by the *new small wheels (NSW)*. Unlike the Run-2 small wheels, the NSW incorporate micromegas (MM) and small-strip TGCs (sTGCs), providing both precision tracking and trigger capabilities in the forward region [108].

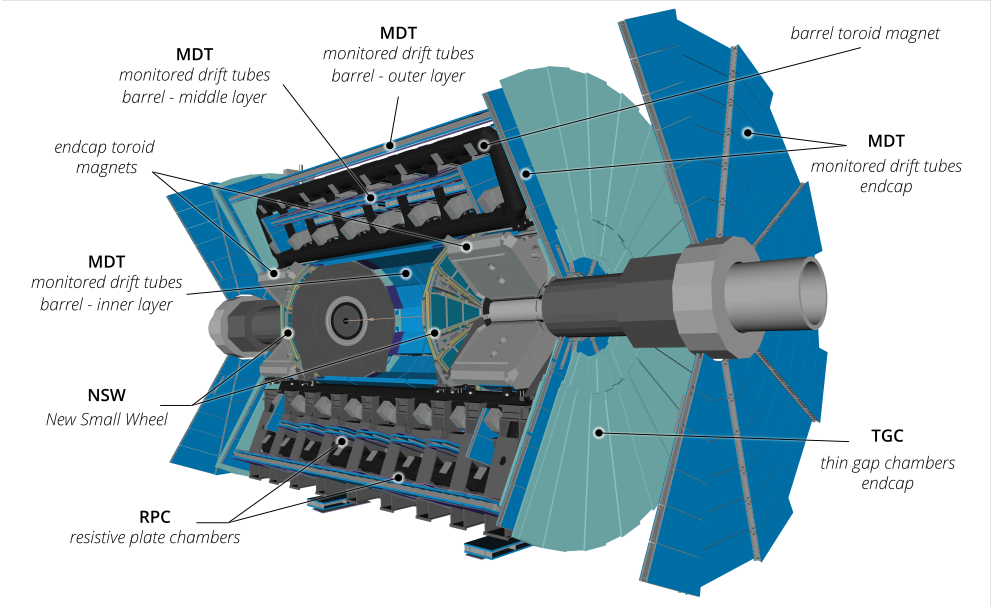


Figure 4.10: Illustration of the ATLAS Muon Spectrometer in Run-3, showing the new small wheels (NSW) that replaced the innermost end-cap wheels from Run-2 [108].

4.2.5 Trigger and data aquisition system

The LHC delivers data at an astounding bunch-crossing rate of 40 MHz. Given current technology, it is impossible to read out and store detector information from every collision. To manage this data rate effectively, the ***ATLAS trigger and data aquisition (TDAQ) system*** applies a set of event selections, known as ***triggers***, to identify only those events relevant for physics analyses. Based on predefined conditions, the system reduces the data rate to approximately 1 kHz, permanently discarding events likely originating from uninteresting scattering processes.

The trigger system consists of two levels. The first is the hardware-based ***Level-1 (L1) trigger***, which uses reduced-granularity information namely from the calorimeters and muon systems to identify high-p_T muons, electrons, photons, jets, τ-leptons, as well as events with large missing transverse energy or large total transverse energy [108]. The L1 trigger reduces the rate to about 100 kHz. Accepted events are read out by detector-specific read-out drivers (RODs) and sent to the readout system (ROS), which is the first common stage of the DAQ system[11]. The ROS buffers the data and forwards it to the second trigger stage, known as the ***high-level trigger (HLT)***. Regions of the detector in which L1 identifies an object[12] in η and ϕ within the detector are defined as regions of interest

[11] For Run-3, the readout path includes new FELIX and SW ROD systems [108].

[12] Here, object refers to any detectable feature, such as calorimeter energy deposits or particle trajectories.

(RoIs). These RoIs are then investigated by the second trigger stage. The HLT is software-based and runs algorithms similar to those used in offline event reconstruction, operating on full-resolution information from all subdetectors, including tracking information from the ID. It further reduces the event rate to about 1 kHz. Events accepted by the HLT are sent to dedicated servers, the subfarm outputs (SFOs), which prepare them for transfer to offline storage, marking the final stage of the DAQ system [108].

Events used in a given physics measurement are selected using an L1 trigger item together with a sequence of HLT algorithms that reconstruct physics objects and apply the required kinematic selections. This sequence of selections is known as a ***trigger chain*** [117]. The full set of chains used during data taking is referred to as the trigger menu, which also specifies the prescales applied to each chain. A ***prescale*** is a value that controls the rate of accepted events: for a prescale value of n, an event has a probability of $1/n$ of being recorded. Prescales can be applied at L1 or at the HLT and may take values greater than or equal to one. The trigger choices used for the measurement presented in this thesis are discussed in Section 7.3.2.

4.2.6 Data collection and analysis

Depending on the trigger chain(s) used, events accepted by the HLT are assigned to ***data streams***, which group events according to their physics signatures or purpose. These streams are then written to separate files for offline processing. The recorded data are organised into distinct periods corresponding to each LHC fill, referred to as Runs, as introduced in Section 4.1.1. Within each run, events are assigned to one or more streams according to the trigger decisions. A run is divided into ***luminosity blocks (LBs)***, which define intervals of approximately constant luminosity and stable detector conditions, typically lasting about one minute. To ensure high-quality data, monitoring algorithms evaluate whether, for a given LB, the detectors are functioning properly, the performance of physics object reconstruction is acceptable, and any data quality defects are tolerable. LBs that fail these criteria are flagged to be ignored in analyses. All LBs from all runs that meet the quality criteria are compiled into a list called the ***Good Runs List (GRL)***, maintained by the ATLAS data quality group, which effectively masks out the bad LBs. Only the GRL LBs are considered suitable for physics analyses, and the integrated luminosity used in a given analysis must account for any excluded data.

4.3 Future prospects and upgrades

Run-3 of the LHC is currently underway, with approximately 351 fb^{-1} of integrated luminosity already delivered. When combined with Run-2 data, the total Run-2+Run-3 dataset corresponds to an impressive ≈ 0.5 ab^{-1} of integrated luminosity[13]. This value is expected to continue increasing steadily through the remainder of Run-3, which is scheduled to conclude in the summer of 2026. Following Run-3, the LHC will enter a long shutdown period (LS3), dedicated to major accelerator upgrades for Run-4, collectively referred to as the ***high-luminosity LHC (HL-LHC)*** [118]. As the name suggests, these upgrades are driven by the goal of significantly extending the discovery potential of the LHC by increasing the collision rate and enabling stable operation over many additional years. The HL-LHC is designed to deliver a centre-of-mass energy of 14 TeV and an order-of-magnitude increase in integrated luminosity, with a target of up to 3000 fb^{-1} accumulated over roughly twelve

[13] The prefix “atto” (a) denotes a factor of 10^{-18}. Thus, 1 ab^{-1} = 10^{18} b^{-1}.

years of operation following the upgrade. Achieving these performance goals requires substantial advances in accelerator technology, including novel superconducting magnet designs, improved beam collimation systems, and superconducting materials with minimal energy dissipation [118].

To cope with the harsher radiation environment and much higher collision rates anticipated at the HL-LHC, the ATLAS detector is undergoing an extensive upgrade, known as the *Phase-II upgrade* [119, 120]. A central component of this upgrade is to preserve the excellent performance of particle reconstruction and identification in the ID, which is critical for nearly all physics analyses. The increased collision rate will lead to significantly higher pile-up, with up to 200 pp interactions per bunch crossing, resulting in extremely high track densities in the innermost detector regions. Therefore, one of the most significant components of the upgrade is the complete replacement of the current inner tracking system with a new, all-silicon detector known as the ***inner tracker (ITk)***. The ITk, illustrated in Figure 4.11, will consist of a high-granularity pixel detector [121] surrounded by a silicon strip detector [122]. The new pixel detector substantially extends the pseudorapidity coverage in the forward region from $|\eta| < 2.5$ to $|\eta| < 4$, increasing the acceptance for key physics processes. In addition, a new high-granularity timing detector (HGTD) [123] will be installed (shown in orange in Figure 4.11) in the forward region to complement the excellent spatial resolution of the ITk, and provide increased radiation hardness alongwith readout electronics capable of handling higher data rates. The Phase-II upgrade is a global effort, with teams from across the world contributing to detector modules, services, and readout systems.

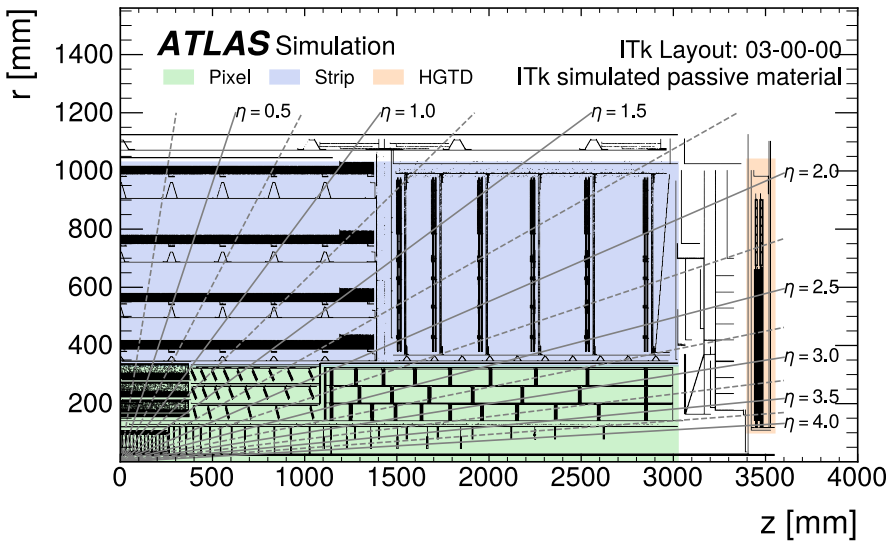


Figure 4.11: An illustration of the future ATLAS ITk, showing the pixel detector (in green), the strip detector (in blue), and the high-granularity timing detector (HGTD, in orange) [119].

By delivering unprecedented integrated luminosity, the HL-LHC is expected to enhance sensitivity to rare SM processes and extend the reach of searches for BSM physics. In this way, the HL-LHC era promises to overcome limitations due to low statistics in certain physics analyses and to significantly broaden the discovery potential of HEP experiments.

CHAPTER 5

Data and simulation of proton-proton collisions in ATLAS

"It is always something insignificant that makes a difference, the way the simple positioning of a comma can change the meaning of an eight-hundred-page novel."

— David Foenkinos, Second Best

Once data from the ATLAS detector is collected, following the procedure described in Section 4.2.6, they are used for physics analyses after undergoing reconstruction and event selections specific to the analysis (Chapters 6-8). As these analyses are guided by theoretical motivations (Chapter 3), the results obtained from data must be compared with corresponding theoretical predictions. Such comparisons cannot rely solely on analytical calculations due to the complexity of pp collisions, which involve hard scattering, parton showering, hadronisation, particle decays, and pile-up, as well as the detailed response of the detector, including effects related to resolution, efficiency, and acceptance. For this reason, ***Monte Carlo (MC) simulation*** techniques are employed. The MC method is a statistical technique that generates possible outcomes of stochastic processes through random sampling from underlying probability distributions. In particle physics, MC techniques are used to simulates collision events while preserving the intrinsic probabilistic nature of the underlying physics processes.

MC samples are generated using dedicated event generators that model the different stages of the physics process and are subsequently passed through a detailed simulation of the ATLAS detector. The events are then reconstructed using the same algorithms as applied to collision data, enabling a consistent and meaningful comparison between prediction and observation. In addition, the underlying simulated particle information is used to characterise the detector response and is often used in the optimisation, calibration, and performance evaluation of reconstruction and identification algorithms (Chapter 6). Discrepancies between data and simulation may indicate limitations in theoretical modelling, imperfections in the detector simulation, or missing physics contributions.

Following the description of the ATLAS dataset in Section 5.1, Section 5.2 outlines the steps involved in the simulation chain and the generators used at each stage. The levels of event description, that are relevant for this thesis, are discussed in Section 5.3, and the specific simulated samples used in this work are detailed in Section 5.4. The theoretical content is based on Refs. [124, 125].

5.1 ATLAS dataset

In this thesis, the full Run-2 dataset collected with the ATLAS detector from pp collisions at $\sqrt{s} = 13$ TeV, corresponding to an integrated luminosity of 140 fb^{-1}, is used. Only data-taking periods satisfying the data-quality requirements for a given analysis are considered. In particular, only the luminosity blocks that belong to the GRL are included (Section 4.2.6).

The raw detector signals are processed and transformed into reconstructed objects such as tracks, jets, and leptons. This procedure is referred to as ***object reconstruction***, and the resulting entities are collectively called *physics objects*. These objects constitute the starting point for the analyses presented in this thesis. In addition to the reconstruction itself, object-specific identification, isolation, and calibrations are applied to ensure that the reconstructed objects are physically meaningful. Further details of these procedures are provided in Chapter 6. These processed physics objects then undergo further analysis-specific event reconstruction and application of selection criteria, described in Chapters 7 and 8.

5.2 Simulation of events

The modelling of each pp collision event is performed in several stages. Revisiting Equation 3.1, the core component of the simulation is the description of the hard interaction, which is achieved through the ***matrix element (ME) calculation***. As discussed in previous chapters, this part of the event is characterised by large momentum transfers over short distance scales, allowing the use of pQCD computed at a fixed order of α_s. In this thesis, event generators capable of performing ME calculations at LO and NLO are employed. The most commonly used generators are POWHEG BOX[1] [126], SHERPA[2] [127], and MADGRAPH5_AMC@NLO[3] [128, 129]. Each of these generators exists in multiple versions, typically denoted as vX.Y or similar numbering schemes, which are specified in Section 5.4.

The ME calculation alone does not provide a complete description of the final state. Additional stages of the event evolution must be taken into account in order to model a realistic pp collision event. An example of the structure of a typical collision event is illustrated in Figure 5.1. Following the hard interaction (shown by the big red filled circle), the outgoing partons can radiate additional gluons, which may subsequently split into quark-antiquark pairs. This iterative process gives rise to a cascade of emissions known as the ***parton shower (PS)*** (shown in blue). Radiation emitted from outgoing partons after the hard scattering is referred to as ***final state radiation (FSR)***, while radiation emitted from incoming partons prior to the hard interaction is known as ***initial state radiation (ISR)***. The two most commonly used generators to model the PS are PYTHIA [130] and HERWIG [131, 132]. In principle, the first and hardest emissions can be described either by higher-order contributions in the ME calculation or by the PS approximation. To prevent double counting of radiation and ensure a consistent description across phase space, a ***ME/PS matching*** procedure is implemented in different generators [124]. The *beam-remnant* partons that do not participate in the hard scattering, and are therefore not included in the ME calculation, can still interact and contribute to the final state. In addition, further interactions between partons from the same pp collision, referred to as

[1] **PO**sitive **W**eight **H**ard **E**mission **G**enerator **BOX**

[2] **S**imulation for **H**igh-**E**nergy **R**eactions of **PA**rticles

[3] MADGRAPH5 **A**utomated **M**ulti-**C**ore @ **N**ext-to-**L**eading **O**rder

multiple-parton interactions (MPIs), may occur. These contributions are collectively referred to as the ***underlying event (UE)*** (shown in purple). The UE is modelled phenomenologically in MC generators, with model parameters adjusted to reproduce experimental data. A specific choice of these parameters within a given generator is referred to as a *tune*.

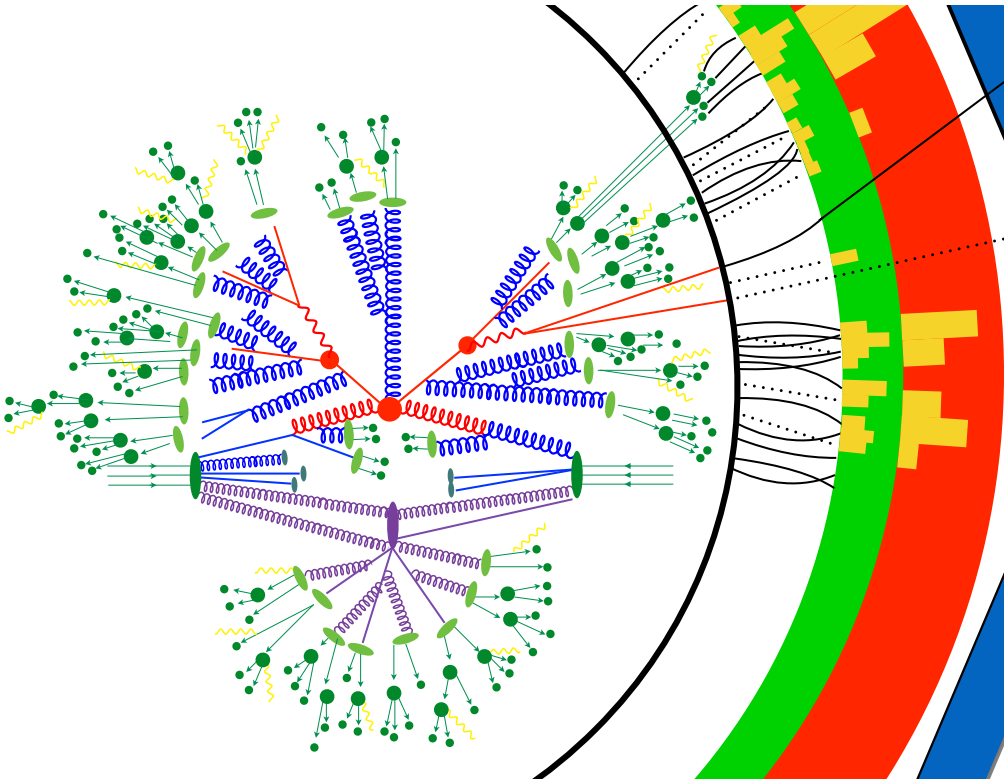


Figure 5.1: Illustration of a pp collision event corresponding to the process $gg \to t\bar{t} \to q'\bar{q}b\,\mu\nu_\mu\bar{b}$, as modelled by an event generator. The incoming partons are depicted as dark-green ovals, while the hard-scattering interaction is represented by the large red filled circle. The produced top quarks are indicated by smaller red filled circles. The subsequent parton shower evolution is shown in blue, and additional soft interactions from the underlying event are illustrated in purple. Final-state partons undergo hadronisation (light-green filled circles), after which the resulting hadrons decay (dark-green filled circles). Photon radiation, shown in yellow, may occur at various stages of the event evolution. This schematic is reproduced from Ref. [133].

The PS originating from both the hard interaction and the UE, evolves towards lower energy scales where α_s becomes large. In this low-perturbative regime, coloured partons cannot exist freely and form hadrons through the process of ***hadronisation*** (shown by light-green filled circles). Since this process occurs at energy scales where pQCD is no longer applicable, it is described phenomenologically. Two main classes of hadronisation models are commonly used in MC event generators: the *Lund string model*, implemented for example in Pythia [130, 134], and the *cluster model*, used in Herwig [131, 132]. The Lund string model, briefly introduced in Section 2.1.2 and illustrated in Figure 2.2, is based on the idea that the colour field between the separating quarks is confined to a narrow flux tube, referred to as a *string*. When sufficient energy is accumulated in the string, it breaks producing a new quark-antiquark pair. This string-breaking process, also known as ***fragmentation***, repeats until the available energy is exhausted, ultimately resulting in the formation of hadrons. The resulting

collimated group of hadrons, originating from a single parton or a small number of nearby partons, is reconstructed experimentally as a hadronic jet (Section 6.4). In the cluster model, gluons first undergo non-perturbative splitting $g \to q\bar{q}$. The resulting colour-connected quark-antiquark pairs combine to form localised colour-singlet objects known as *clusters*, which differ from strings in that they are spatially localised rather than extended objects. If a cluster has a mass above a certain threshold, it undergoes further splitting into lighter clusters. Each cluster then decays, typically isotropically, into hadrons.

As the final step of event generation, the decays of unstable hadrons are simulated (shown by small dark-green filled circles). A particle is typically considered *stable* if it has a lifetime greater than approximately 30 ps, corresponding to particles that typically do not decay before reaching the first layer of the detector. Following this decay stage, the simulation chain proceeds to the ATLAS detector simulation, where the detector response is modelled. This can be performed via *full simulation*, based on GEANT4 [135], which provides a detailed description of particle interactions with the detector material, or using a *fast simulation* approach. In the latter, the calorimeter response is replaced by a parametric description [136][4]. The effect of pile-up is accounted for through a procedure known as ***pile-up overlay***, in which events from additional pp collisions, typically from soft QCD processes, are overlaid onto the simulated hard-scatter event. In this work, these additional interactions are simulated using Pythia 8 [130] and using the A3 set of tuned parameters [137]. The resulting events are subsequently reweighted to ensure that the distribution of the average number of interactions per bunch crossing matches that observed in the recorded data. Finally, the simulated events are passed through the ***digitisation*** process, in which the energy deposits and particle interactions in the detector are converted into electronic signals, reproducing the detector readout and yielding output in a format equivalent to that of the real data from the DAQ system.

5.3 Levels of event description

Simulating a pp collision in multiple stages offers the advantage of accessing an event at different points in its development. Depending on the analysis being performed, these distinct levels of event description enable the extraction of meaningful results. In the entanglement studies presented in this thesis, three levels are particularly relevant: the detector level, obtained after the final simulation step described in Section 5.2, as well as the parton and particle levels. The definitions of these levels and their respective roles in this work are summarised below.

5.3.1 Detector level

The ***detector level*** refers to the stage where simulated events are obtained after the full detector simulation based on GEANT4. The resulting simulated detector signals are processed using the same reconstruction algorithms as applied to collision data, producing reconstructed physics objects (Chapter 6). Unless otherwise stated, these objects undergo the same identification, isolation, calibration, and analysis-specific selection requirements as applied to data.

[4] Although fast simulation significantly reduces the CPU time required for event generation, it does not reproduce jet substructure observables as accurately as the full simulation. Therefore, fast-simulation samples are primarily used for the estimation of modelling uncertainties (Section 7.9).

In the following chapters, the terms ***detector level*** and ***reconstruction level*** are used interchangeably to denote events after detector simulation and object reconstruction, that is the stage equivalent to reconstructed collision data.

5.3.2 Parton level

The ***parton level*** refers to the stage of the event record that contains the quarks and gluons produced in the hard scattering process, prior to hadronisation. This level includes objects both before and after FSR, but before the full PS. At this level, the decay products of the top quark are represented by partons[5]. In particular, the quarks originating from the hadronically decaying top quark have not yet undergone hadronisation into jets. No detector simulation or reconstruction is applied at this stage.

This level is especially relevant for the entanglement measurement, since the ***entanglement condition***, defined in Equation 3.12, ***is valid only at the parton level***. This stage therefore provides direct access to the true value of the observable $D(D_3)$, as discussed in Chapter 7(8).

5.3.3 Particle level

The ***particle level*** corresponds to the stage at which physics objects are defined using stable particles before any detector simulation is applied. Particle-level objects, such as electrons, muons, and jets, are constructed to closely correspond to the detector-level objects used in the measurement (Section 7.1). The objects at this level are also sometimes referred to as *truth-level objects*.

While both the parton and detector levels can be used to derive meaningful values of the entanglement observable, the particle level plays a particularly important role in the threshold entanglement measurement. This aspect is discussed in detail in Chapter 7.

A pictorial illustration of the three levels of event description discussed here, shown for the lepton+jets resolved decay channel, is provided in Figure 5.2.

5.4 Simulated samples

Among the simulated samples, the ***signal samples*** describe the targeted physics process, while the ***background samples*** correspond to processes that either mimic the signal or lead to the same final state. For a given final state, here the $t\bar{t}$ lepton+jets channel, all such signal and background contributions must be considered.

For the entanglement measurement at threshold, the signal consists of SM $t\bar{t}$ production described in pQCD, as well as the NRQCD $t\bar{t}$ simulations at threshold that also include the formation of colour-singlet quasi-bound states. The simplified pseudo-scalar resonance model is also considered for comparison. The background contributions arise from single-top production; W+jets production, where the W boson decays leptonically; Z+jets production, with the Z boson decaying leptonically; $t\bar{t}$ production in association with a SM Higgs boson ($t\bar{t}H$) or with a vector boson ($t\bar{t}V$, where $V = W, Z$); and diboson production (ZZ, WZ, WW). An additional background originates from physics objects that are misidentified as *prompt leptons*, i.e as leptons that originate from the hard-scattering process (Section 6.1). This background arises from two sources: i) leptons arising from semi-leptonic decays

[5] Note that for the parton level, leptons are also referred to as "partons", which is different from the usual QCD definition of the word.

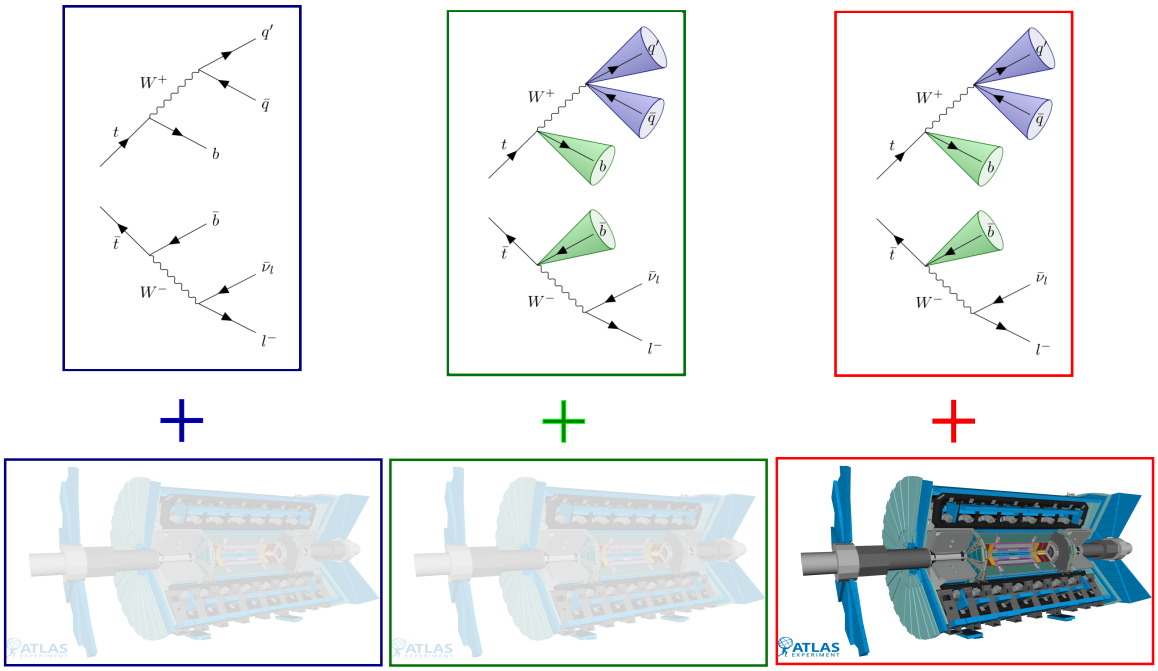


Figure 5.2: Illustration of the different levels of event description for the $t\bar{t}$ lepton+jets resolved decay channel: parton level (left, blue), particle level (middle, green), and detector level (right, red). At the parton level, the decay products are represented by quarks and leptons before hadronisation. At the particle and detector levels, the quarks have hadronised into particles that are clustered into jets, illustrated by the conical jet shapes in the diagram. The faded ATLAS detector in the first two panels indicates that detector effects are not included at the parton or particle levels.

of hadrons or from photon conversions in the detector material, referred to as *non-prompt leptons*, and ii) objects reconstructed as leptons that do not correspond to real leptons, referred to as *fake leptons*. Collectively, this background contribution is called non-prompt/fake lepton background. The diversity of processes by which these non-prompt and fake leptons are produced, as well their strong dependence on instrumental effects, makes this background difficult to model reliably in simulation. Therefore, their contribution is estimated using data-driven techniques. The method employed for this estimation is described in Section 7.8.

For the boosted entanglement measurement, the selected region is expected to be dominated by the SM $t\bar{t}$ signal. The backgrounds are the same as in the threshold analysis, with the non-prompt/fake lepton background evaluated using a method specific to the boosted topology analysis[6].

Except for samples generated with SHERPA, the decays of bottom and charm hadrons are simulated using the EVTGEN program [138]. For all samples described below that are interfaced with PYTHIA

[6] At the time of writing, the non-prompt/fake lepton background estimation was still under study. Therefore, the final boosted entanglement results presented in this thesis do not include this background contribution.

8 [130] for PS and hadronisation, the A14 tune [139] is used together with the NNPDF2.3$_{\text{LO}}$ PDFs set [140]. Each process has three different simulated samples for the data-taking periods of 2015+2016, 2017, and 2018. This is done to account for the changing detector and pile-up conditions across the run periods. Unless explicitly stated otherwise, simulated samples for a given process described below are produced using the same generator versions in both the threshold and boosted entanglement measurements.

5.4.1 SM $t\bar{t}$ production

The SM $t\bar{t}$ production is modelled using the Powheg Box v2 generator [126, 141–143], which provides the ME calculation at NLO in α_s, together with the NNPDF3.0$_{\text{NLO}}$ PDF [144]. The top-quark mass is set to $m_{\text{top}} = 172.5\,\text{GeV}$. The renormalisation and factorisation scales (μ_R and μ_F) are chosen according to the default functional form $\mu_R = \mu_F = \sqrt{m_{\text{top}}^2 + p_{\text{T}}^2}$. The Powheg parameter h_{damp}, which controls the ME/PS matching and effectively regulates ("dampens") the amount of high-p_T radiation, is set to $1.5m_{\text{top}}$ [145]. The generated events are interfaced to Pythia 8.230 for PS and hadronisation. Several alternative MC samples are used to evaluate uncertainties associated with the modelling of the $t\bar{t}$ process. These samples are described in Sections 7.9 and 8.7.2.

The nominal $t\bar{t}$ sample and all alternative samples are normalised to the inclusive $t\bar{t}$ production cross-section calculated at NNLO in QCD, including the resummation of next-to-next-to-leading-logarithmic (NNLL) soft-gluon calculated using Top++ 2.0 [146–152]. For $\sqrt{s} = 13\,\text{TeV}$ and $m_{\text{top}} = 172.5\,\text{GeV}$, the calculated cross-section is $\sigma(t\bar{t})_{\text{NNLO+NNLL}} = 832 \pm 51\,\text{pb}$. For the threshold entanglement measurement, the kinematic variables of the $t\bar{t}$ system in the nominal and all alternative pQCD $t\bar{t}$ samples are corrected to match more precise predictions computed at NNLO in QCD and NLO in electroweak (EW). These corrections are derived following the method established in recent measurement of $t\bar{t}$ production at threshold analysis [21]. In this approach, Matrix programme [153] is used to obtain the NNLO QCD predictions in α_s, while Hathor [154–156] provides the NLO EW predictions. In both programs, μ_R and μ_F are set to the same values as in the ME generator. The resulting predictions are used to derive a two-dimensional, event-by-event kinematic reweighting in the variables $(m_{t\bar{t}}, \cos\theta^*)$, where θ^* is defined as the angle between the momentum of the top quark in the $t\bar{t}$ centre-of-mass frame and the momentum of the reconstructed $t\bar{t}$ system in the laboratory frame. The corrections are provided in the ranges $340 \leq m_{t\bar{t}} \leq 3500\,\text{GeV}$ and $-1 \leq \cos\theta^* \leq +1$. Since both the NNLO-QCD and NLO-EW predictions assume on-shell top quarks, no corrections are applied to events with $m_{t\bar{t}} < 340\,\text{GeV}$. Since this procedure is implemented through an event-by-event reweighting, it is referred to as ***NNLO reweighting***. The associated systematic uncertainties of this method are discussed in Section 7.9. This reweighting is applied exclusively in the threshold entanglement measurement[7].

5.4.2 Single-top production

At the LHC, single top-quark production proceeds via three electroweak processes at LO, commonly referred to as the *canonical channels*:

[7] While the reweighting would also be beneficial for the boosted analysis, a dedicated reweighting procedure had not been implemented within the timeframe of this thesis.

- ***s*-channel**: A quark–antiquark annihilation process mediated by the exchange of a time-like virtual W boson.

- ***t*-channel**: A process in which a top quark is produced via the exchange of a space-like virtual W boson.

- ***Wt* associated production**: Production of a top quark in association with an on-shell W boson.

Among these processes, Wt associated production constitutes the dominant source of single-top events in this work. This process is simulated at NLO in QCD using Powheg Box v2 [126, 142, 143] with the NNPDF3.0$_{\text{NLO}}$ [144] PDF set. Beyond LO, several Feynman diagrams contributing to the Wt process are shown in Figure 5.3. These diagrams can be interpreted as LO $t\bar{t}$ production with the $\bar{t}$ subsequently decaying into $\bar{b}W$. This leads to an overlap between the Wt and $t\bar{t}$ processes. To treat this overlap and avoid double counting, the nominal Wt sample is generated using the *diagram removal (DR)* scheme [145, 157], in which the doubly resonant $t\bar{t}$-like contributions are explicitly removed at the amplitude level. The uncertainty associated with this treatment is evaluated by comparing the nominal DR sample with an alternative sample generated using the *diagram subtraction (DS)* scheme where these contributions are removed at cross-section level [158], as described in Section 7.9.4.

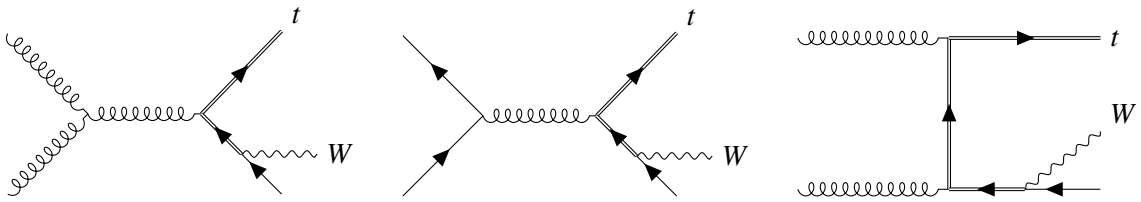


Figure 5.3: Feynman diagrams of single-top production in the Wt channel beyond LO. Top and antitop quarks are represented with double lines.

The s- and t-channel processes are generated using the same setup as the Wt sample. For the t-channel process, top-quark decays are handled at LO using Madspin [159, 160] in order to preserve spin correlations. All three production modes are interfaced with Pythia 8.230 for PS, hadronisation, and UE modelling. The simulated samples are normalised to their respective theoretical cross-sections: NNLO+NNLL accuracy for the Wt process [161], and NLO accuracy for the s- and t-channel processes [162, 163].

5.4.3 W+jets and Z+jets

The production of W and Z bosons in association with jets is simulated using Sherpa 2.2.11 and Sherpa 2.2.14 [164]. The ME is computed at NLO for final states with up to two partons, and at LO accuracy for configurations with up to four partons. The calculations are performed using the Comix [165] and OpenLoops [166] libraries.

The MEs are matched with the Sherpa PS [167] using the MEPS@NLO prescription [168–171], employing the parameter tune recommended by the Sherpa authors. The NNPDF3.0$_{\text{NNLO}}$ PDF set is used and the generated samples are normalised to an NNLO prediction [172].

5.4.4 $t\bar{t}V$ and $t\bar{t}H$

Background processes involving the production of a top-quark pair in association with a vector boson ($t\bar{t}V$, with $V = W, Z$) are simulated at NLO accuracy in QCD using MadGraph5_aMC@NLO 2.3.3 [128] generator with the NNPDF3.0$_{\text{NLO}}$ PDF set [144]. The generated events are interfaced with Pythia 8.210 for PS, hadronisation, and UE modelling. For the boosted entanglement measurement, the $t\bar{t}Z$ process is instead generated using the newer MadGraph5_aMC@NLO 3.5.1 and interfaced with Pythia 8.309. All samples are normalised to theoretical cross-sections calculated at NLO accuracy in both QCD and EW [173].

The ME generation of $t\bar{t}H$ events is performed with Powheg Box v2 [126, 141–143, 174] at NLO using the NNPDF3.0$_{\text{NLO}}$ PDF set [144]. The events are interfaced with Pythia 8.230 and are normalised to the corresponding predicted cross-section [173].

5.4.5 Diboson

Samples of diboson (VV, where $VV \in \{WW, WZ, ZZ\}$) processes are generated using Sherpa 2.2.1 or Sherpa 2.2.2 [127], depending on the specific final state. For the boosted entanglement measurement, the samples are instead generated using Sherpa 2.2.14. The simulation includes off-shell effects and, where relevant, Higgs boson contributions. Fully leptonic and semi-leptonic final states are simulated using MEs calculated at NLO accuracy in QCD for configurations with up to one additional parton, and at LO accuracy for up to three additional parton emissions. Loop-induced $gg \to VV$ processes are generated separately using LO-accurate MEs, including up to one additional parton emission, for both fully leptonic and semi-leptonic final states. The ME calculations are matched and merged with the Sherpa PS using the MEPS@NLO prescription [168–171], based on Catani–Seymour dipole factorisation [175, 176]. Virtual QCD corrections are provided by the OpenLoops library [166, 177, 178]. The NNPDF3.0$_{\text{NNLO}}$ PDF set [144] is used throughout, together with the dedicated PS tune provided by the Sherpa authors.

5.4.6 $t\bar{t}$ quasi-bound states

Near the $t\bar{t}$ production threshold, where NRQCD effects becomes relevant, the dominant contribution arises from colour-singlet $t\bar{t}$ production via gg fusion. In this regime, the attractive QCD potential between the top and antitop quarks can lead to the formation of quasi-bound-state effects, as described in Section 3.3. In the threshold entanglement analysis, these NRQCD effects are modelled using the approach described in Ref. [95]. A Green's function in the Coulomb gauge of the NRQCD framework is introduced to describe the momentum distribution of top quarks in the presence of the QCD potential. This function is computed numerically by solving the Lippmann-Schwinger equation, using as input the Fourier transform of the QCD potential [179, 180]. It is then used to reweight LO pQCD MEs for the production of an (off-shell) colour-singlet $t\bar{t}$ pair in an S-wave configuration near threshold. This reweighting enables the simulation of quasi-bound-state $t\bar{t}$ events and effectively includes the full spectrum of $^1S_0$ mass states. Only gg fusion initiated processes, which contribute to the colour-singlet $^1S_0$ states, are included, while subleading $q\bar{q}$-initiated contributions such as $^3S_1$ states, are not considered in this model.

The full production and decay chain into a six-body final state is generated using MadGraph 3.4.2 with the NNPDF3.0$_{\text{NLO}}$ PDF set, and interfaced with Pythia 8.3 for PS and hadronisation. To remain consistent with the calculations presented in Ref. [95], the top-quark mass and width are set to 173 GeV

and 1.49 GeV, respectively. The reweighting is applied only to events with $m_{t\bar{t}} < 350$ GeV and with top-quark momentum magnitude $p* < 50$ GeV in the $t\bar{t}$ rest frame, evaluated at the parton level. This selection ensures that the kinematics that are included are non-relativistic. In order to jointly study this NRQCD-inspired sample together with the SM pQCD $t\bar{t}$ production sample (Section 5.4.1), which otherwise does not include quasi-bound-state effects, care must be taken to avoid double counting of contributions that are present in both the pQCD prediction and the dedicated threshold description. In particular, the pQCD $t\bar{t}$ sample may already incorporate NRQCD threshold effects up to NLO accuracy. A double-counting correction procedure has been proposed in Ref. [95], but is not applied here, under the assumption that the overlap is minimal. The resulting NRQCD sample is normalised to a total cross-section of 6.43 pb, obtained analytically including all decay channels [179]. No uncertainty estimate is currently available for this value. This sample corresponds to the nominal $t\bar{t}$ quasi-bound-state model used for the threshold entanglement measurement presented in this thesis. It provides a more complete description of $t\bar{t}$ quasi-bound-state dynamics through ***Green's functions reweighting (GFRW)***. In line with the name of the reweighting procedure, it is therefore denoted $\boldsymbol{t\bar{t}_{\mathrm{GFRW}}}$ throughout this thesis.

For comparison, the simplified pseudo-scalar resonance model [94] is also considered. In contrast to the nominal $t\bar{t}_{\mathrm{GFRW}}$ model, this approach approximates the threshold quasi-bound-state dynamics using a single pseudo-scalar resonance corresponding to the lightest colour-singlet ${}^1S_0^{[1]}$ configuration, and is denoted by $\boldsymbol{\eta_t}$. Event generation for this sample is performed with MadGraph 3.5.5 using the NNPDF3.0$_{\mathrm{NLO}}$ PDF set. The full production and decay chain to the four-body final state ($WbWb$) is included at ME level, while PS and hadronisation are simulated using Pythia 8.3. The η_t mass is set to 343 GeV, with a total width of 2.8 GeV. This choice reflects the expectation that the mass of the quasi-bound state is approximately $2 \cdot m_{\mathrm{top}}$ minus a binding energy of about 2 GeV [90, 179]. The width is taken to be roughly twice that of the top quark which is appropriate for the ground-state of the quasi-bound system [19]. No kinematic requirements are imposed on $m_{t\bar{t}}$ or $p*$. As in the $t\bar{t}_{\mathrm{GFRW}}$ sample, this sample is also normalised to a total cross-section of 6.43 pb.

CHAPTER 6

Reconstruction and identification of physics objects in the ATLAS detector

"You can touch everything and be connected to nothing."

— Mitch Albom, Have a Little Faith

The measurement of entanglement in top-quark pairs targets final states involving leptons as well as quarks that hadronise into jets. A precise physics interpretation therefore requires the reconstruction and identification of these objects from the raw detector signals recorded in data, as well as from the simulated detector response described in Chapter 5.

Figure 6.1 illustrates the paths of different particles and their interactions with the various components of the ATLAS detector as they traverse through it. Photons and electrons produce electromagnetic (EM) showers in the electromagnetic calorimeter, with little to no energy deposited in the hadronic calorimeter. Muons, in contrast, traverse most of the detector material with minimal interaction and are typically only detected in the MS and the ID. Neutrinos leave no signature in the detector and therefore escape detection. Other particles, such as protons and neutrons, predominantly deposit energy in the hadronic calorimeters due to strong interaction between the hadrons and atomic nuclei. Additionally, electrically charged particles show bent trajectories as a result of the Lorentz force [181] acting on them in the magnetic field of the detector. The direction of curvature depends on the charge of the particle and allows for the determination of its momentum and charge. The signals produced by these interactions are recorded as raw detector data. These signatures are subsequently used by reconstruction algorithms and analysis tools to identify the particles and measure their properties for use in physics analyses.

In ATLAS, this reconstruction is performed within the central ***Athena*** software framework [182], where information from all subdetectors is combined to form high-level objects such as electrons, muons, jets, and missing transverse momentum. Data formats containing reconstructed and often calibrated objects are then provided to physics analyses, where they are used to define specialised event selections and reconstruct composite objects such as top quarks or Higgs bosons.

The performance of the algorithms used to reconstruct these objects are overseen by dedicated *combined performance (CP)* groups within the ATLAS collaboration. Each CP group is responsible for a specific class of objects. Differences in the performance of reconstruction algorithms between

data and simulation are accounted for by calibrations derived by these groups. The CP groups periodically provide official *recommendations*, which specify the necessary analysis settings such as calibrations, identification criteria, and associated systematic uncertainties. These recommendations evolve with time as improved detector calibrations, refined pile-up treatments, or updated algorithms such as modern flavour-tagging methods, become available.

When substantial improvements in detector modelling, reconstruction or tagging algorithms are introduced, ATLAS may perform a full reprocessing of the recorded collision data, which are then referred to as *reprocessed data*. In such a procedure, the raw detector data is reconstructed again using an updated version of the reconstruction software, referred to as a new *release*, producing new reconstructed datasets and corresponding derivations. To maintain consistency, simulated MC samples are reconstructed with the same software release as the reprocessed data.

The entanglement analyses considered in this work rely on different software releases and analysis frameworks within the ATLAS collaboration software environment. The threshold entanglement analysis is based on Athena release 21, where CP recommendations are implemented through the *AnalysisTop* framework, which is a dedicated top-quark analysis framework developed for Run-2. In contrast, the boosted entanglement analysis, being a more recent study, is based on Athena release 25 using the modern *TopCPToolkit* framework. TopCPToolkit is the Run-3 successor to AnalysisTop and is built around the common CP tools and analysis algorithms available in Athena release 25. Due to structural changes in the software architecture between releases, certain newer algorithms and calibration schemes are not backward compatible with release 21. Therefore, the boosted analysis benefits from improved reconstruction algorithms, updated calibrations, and modern flavour-tagging techniques available in release 25, whereas the threshold analysis remains tied to the software environment of release 21.

This chapter details the procedures used for the reconstruction, identification, isolation, calibration, and tagging of the relevant physics objects. Unless explicitly stated otherwise, the object definitions and methodologies are consistent between the two analyses, with differences arising only from release-dependent recommendations and software infrastructure. The exact selection criteria applied to these objects are presented in the chapters dedicated to the analyses (Chapters 7 and 8).

6.1 Charged-particle track reconstruction

When charged particles pass through the silicon detectors of the ATLAS ID, they deposit energy in the ID sensors primarily via ionisation and radiative processes. If the energy deposited in a given detector channel (i.e the smallest readout unit of the sensor, corresponding to a single pixel or strip) exceeds a predefined threshold, the signal is recorded as a ***hit***. These pre-threshold values are chosen to ensure high particle detection efficiency while suppressing electronic noise and other instrumental effects that could otherwise be misidentified as a true signal. Neighbouring channels with hits are subsequently grouped together using a connected component analysis approach [26], creating localised energy deposits referred to as ***clusters***. These clusters indicate the presence of charged particles within the sensors and provide the basis for reconstructing their trajectories, referred to as ***tracks***.

Each reconstructed track is characterised by five parameters, defined with respect to a chosen reference point. These parameters are illustrated in Figure 6.2. The ***polar angle*** of the track in the longitudinal plane is denoted by θ, and the ***azimuthal angle*** in the transverse plane by ϕ. The ***transverse impact parameter***, d_0, is the distance of closest approach between the particle trajectory

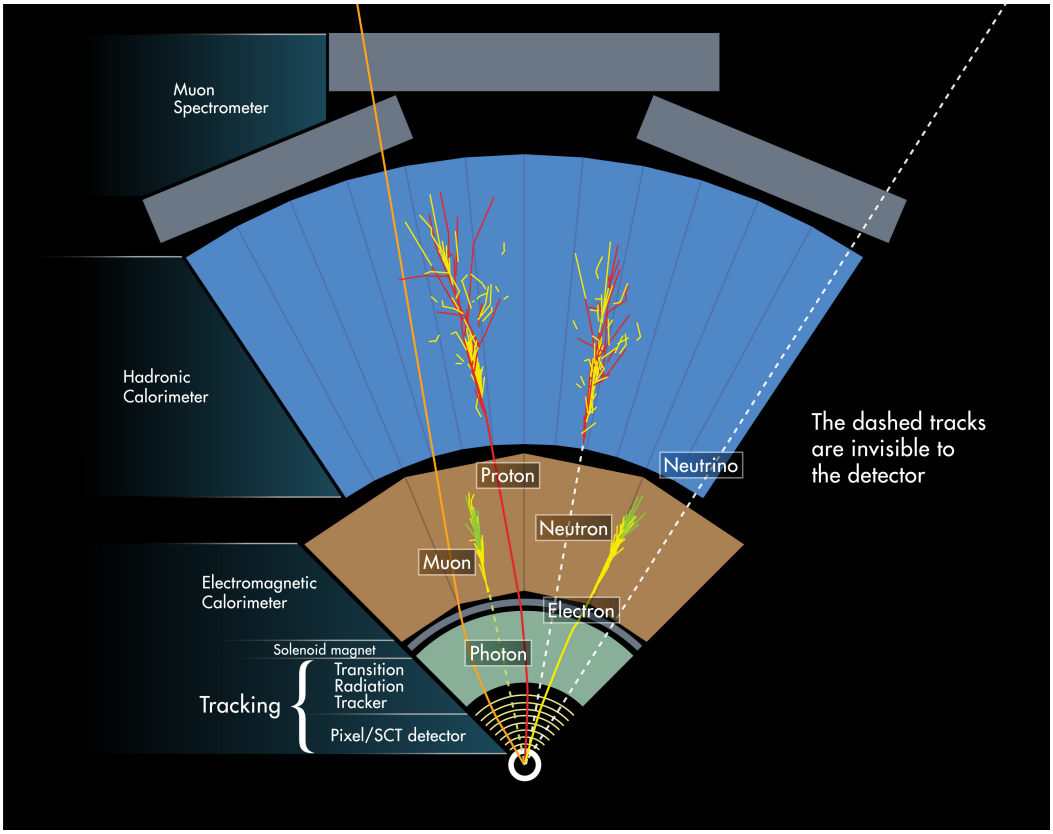


Figure 6.1: Schematic illustration of a sector of the cross-section of the ATLAS detector, showing the trajectories of different particles as they traverse the various detector components. Dashed lines indicate particles that escape detection. Figure adapted from Ref. [183].

and the reference point in the transverse plane. The ***longitudinal impact parameter***, z_0, is defined as the minimum distance between the particle trajectory and the reference point, in the longitudinal plane. The curvature of the track in the magnetic field is quantified by the ratio of the particle charge to its momentum, q/p. The reference line used by default is the beamline, defined as the line that best represents the average position and direction of pp collisions, referred to as the ***beam spot*** or *luminous region* [184]. The corresponding reference point is typically taken to be the centre of this region. However, it is also common to instead use the reconstructed PV (Section 4.2.1), which corresponds to the location of the hardest pp interaction in the event [185].

The reconstruction of charged-particle tracks, or tracking, proceeds through a sequence of steps that combine information from clusters in the pixel detector, SCT, and TRT. A more detailed description of each step can be found in Ref. [26].

Clusterisation As charged particles traverse through the pixel and SCT layers, they typically exhibit an angular trajectory, resulting in the activation of several adjacent sensors within a layer. The clustering procedure is used to group pixels and strips within a given sensor into clusters, where adjacent channels with charge deposits above a given threshold and sharing a common edge or corner are combined. When the charged deposits originate from a single particle, the resulting cluster is

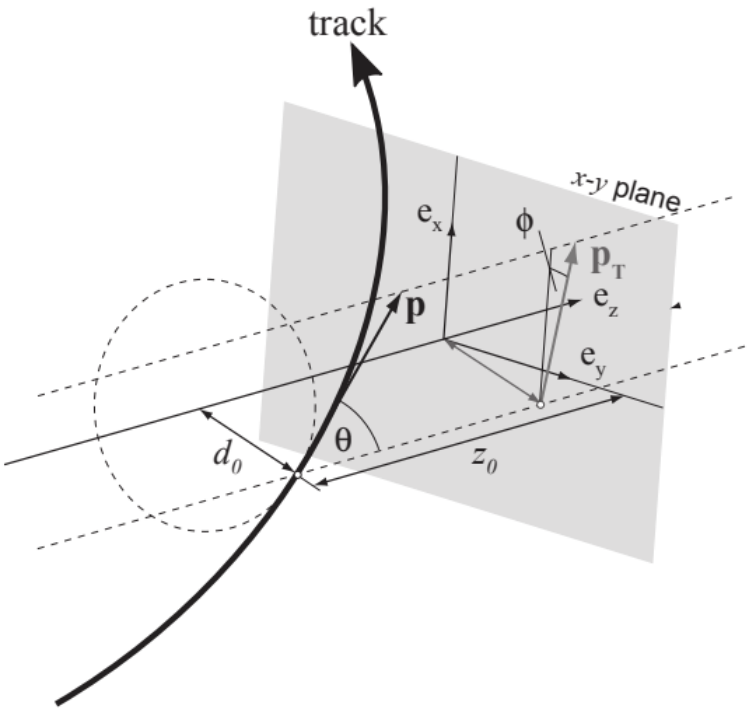


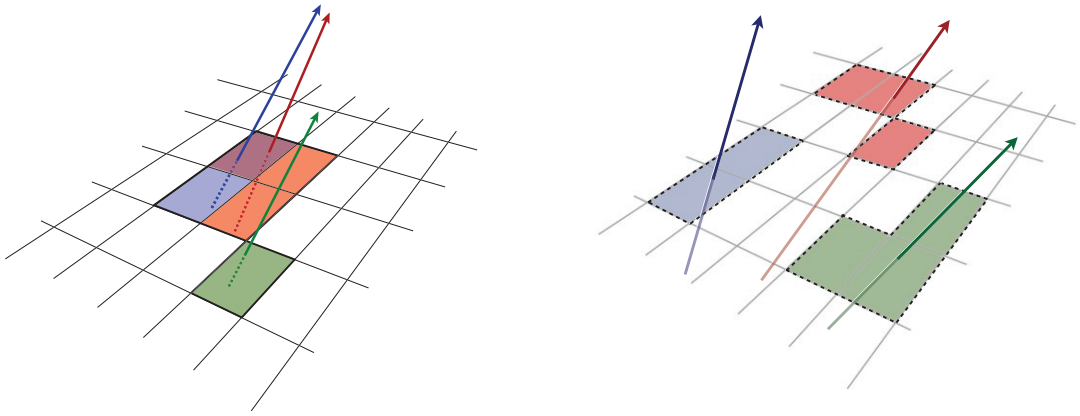
Figure 6.2: Schematic illustration of the perigee representation of the track parameters, where the perigee corresponds to the point of closest approach of the trajectory to the reference point [186].

referred to as a *single particle cluster*. An example of such a cluster is shown in Figure 6.3(a).

In dense environments, characterised by highly collimated charged particles, clusters originating from different particles can no longer be resolved as separate clusters. In such cases, the spatial separation between charged particles becomes comparable to the finite pixel pitch, leading to overlapping charge deposits. This leads to the formation of *merged clusters*, where charge deposits from multiple particles are combined, as illustrated in Figure 6.3(b). These cluster merging effects are particularly relevant for the analyses presented in this thesis, which involve high-energy jets with a large multiplicity of collimated particles. Accurate reconstruction of tracks in these environments is therefore essential for algorithms used to identify the flavour of the quark from which the jet originates. These algorithms rely heavily on precise tracking and vertex reconstruction (Section 6.4.6).

Following cluster formation, the resulting clusters provide three-dimensional measurements known as *space-points*. These points mark the trajectory of the particle through the ID. In the pixel detector, each cluster corresponds directly to a space-point. In the SCT, on the other hand, a space-point is reconstructed by combining clusters from both sides of a strip layer.

Track finding and ambiguity solver Track reconstruction begins with the formation of *track seeds*, built by grouping space-points in sets of three from adjacent or neighbouring layers. These seeds provide initial estimates of possible tracks and their parameters, as shown in blue in Figure 6.4. Since there are typically more track seeds than expected tracks, the inconsistent seed combinations are rejected using criteria based on detector geometry, event topology, momentum, and impact parameters.



(a) Single-particle pixel clusters

(b) Merged pixel cluster

Figure 6.3: Different cluster topologies formed on a pixel sensor of a pixel detector layer. Arrows indicate particle tracks, and colours correspond to energy deposits from individual particles [187].

The chosen seeds are extended into track candidates using a *combinatorial Kalman filter* (CKF) [188], which incorporates further space-points from the remaining layers of the ID while maintaining compatibility with the initial trajectory.

In the next stage of track reconstruction, track candidates are passed to the ***ambiguity solver*** algorithm, which selects the most promising candidates and resolves cases where clusters are assigned to multiple candidates. This stage is particularly important for tracking performance in dense environments, where the probability of incorrect combinations of unrelated clusters increases. This stage also reduces computational cost by selecting the most relevant tracks prior to the final track fitting stage, which is computationally intensive.

Track candidates are assigned scores that favour well-measured, energetic tracks and penalise candidates with holes[1] or incorrectly assigned clusters. Candidates that fail basic quality requirements on the transverse (d_0) and longitudinal (z_0) impact parameters are discarded.

After assigning track scores, the ambiguity solver resolves cases where clusters are shared between multiple track candidates. Clusters are not allowed to be shared by more than two tracks unless they are identified as merged, which means they originate from multiple charged particles. *Shared clusters*, that are not identified as merged, may originate either from falsely reconstructed tracks, called *fake tracks*, or because the cluster has indeed been activated by more than one ionising particle.

The identification of merged clusters is performed using deep neural networks [189], which are trained on features such as deposited energy, pixel sensor positions, and incidence angles of the particles [189]. This aids the ambiguity solver in penalising tracks sharing a cluster which is not identified as merged by the neural networks. Preference is given to tracks with higher track scores when assigning a shared cluster to a track. In addition, a track candidate is not allowed to contain more than two shared clusters. The candidates must satisfy additional quality requirements [26] in order to proceed to the next reconstruction stage.

While shared clusters most likely indicate fake tracks, in dense environments true charged-particle trajectories may also share clusters due to the high collimation of tracks. As a result, the requirements

[1] Holes represent regions in the detector element intersected by a track candidate without creating a matching cluster.

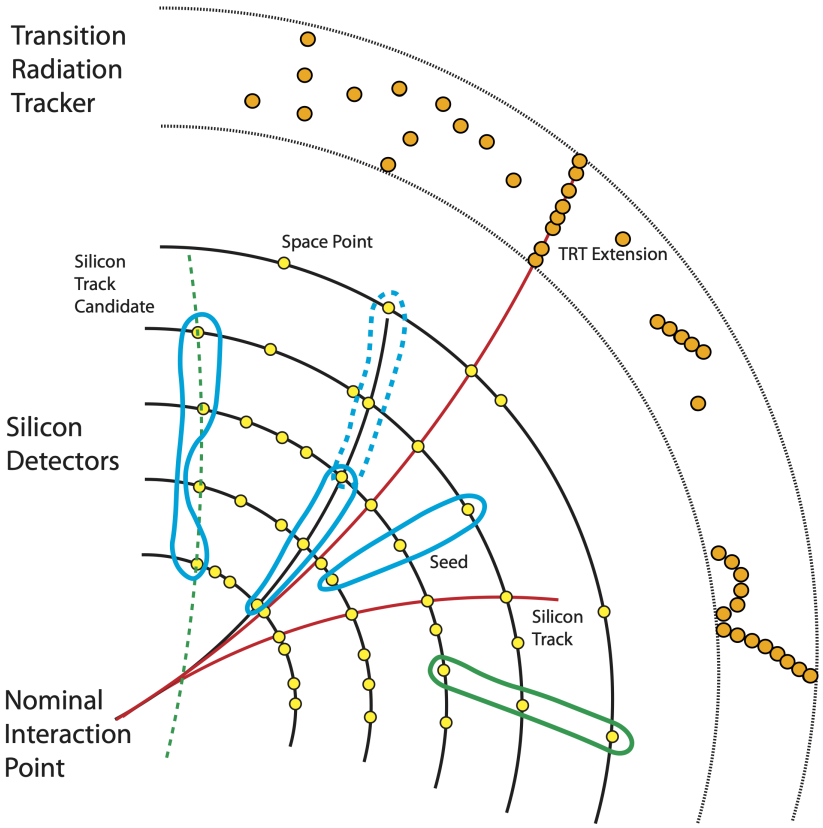


Figure 6.4: Illustration of the ATLAS ID track reconstruction procedure. Space points are shown in yellow, while seeds formed from combinations of space points are indicated in blue. The dashed blue seed illustrates a situation where two seeds correspond to the trajectory of the same charged particle. A seed rejected based on the requirement that it needs to be consistent with a particle originating from the nominal interaction point is shown in green. The black track represents a track candidate reconstructed using CKF. The red tracks correspond to the fully reconstructed silicon tracks obtained after the global χ^2 fit, where an extension into the TRT may also be performed [26].

imposed by the ambiguity solver on the number of shared clusters can lead to genuine tracks being incorrectly rejected as fake tracks, thereby reducing tracking efficiency. Proper identification of merged clusters is therefore essential to suppress fake tracks while preserving the efficiency for reconstructing genuine charged-particle trajectories.

Track fit Track candidates that survive ambiguity resolution undergo a high-resolution χ^2 *fit* to determine their final parameters. For this step, the positions and uncertainties of clusters identified as merged are determined using a *Mixed Density Network (MDN)* [190], which allows accurate parameter determination even when clusters are shared by multiple tracks. Finally, tracks are extended into the TRT where candidate TRT hits are incorporated using another Kalman filter. If the extension is successful, the full track is refitted with the global χ^2 fitter [191]. Incorporating TRT hits improves momentum resolution and provides additional information for particle identification.

The charged-particle track reconstruction described above is commonly referred to as the *inside-out algorithm*. Particles produced at or very close to the PV are referred to as ***prompt particles***. However, many interesting physics signatures involve ***non-prompt particles***, which originate away from the PV. Examples include electrons originating from photon conversions in the detector material, or displaced decay products of long-lived particles such as kaons or B-hadrons, which give rise to secondary vertices (Section 6.4.6). In some cases, cascade decays of heavy-flavor hadrons can also produce tertiary vertices, even further displaced from the pp interaction point. To reconstruct such displaced trajectories, an *outside-in algorithm* is employed [192]. Unlike the inside-out reconstruction, which seeds tracks from space-points in the silicon detectors, the outside-in approach begins with hit segments in the TRT, which are then extended back into the silicon detectors to recover displaced tracks. This approach enhances the efficiency of reconstructing non-prompt particles from secondary and tertiary vertices by utilising detector hits that were not assigned to any tracks in the inside-out reconstruction.

In dense environments, however, the track reconstruction efficiency can be significantly reduced. Due to the increased probability of cluster merging, limitations in the merged-cluster identification and the requirements on shared clusters imposed by the ambiguity solver can lead to some charged-particle tracks failing to be reconstructed, resulting in a loss of tracks during reconstruction. This effect is particularly relevant for physics analyses involving highly energetic jets, such as those studied in this thesis. Jets with high transverse momentum, p_T, typically contain a large number of collimated charged particles, creating a dense tracking environment in the ID. Accurate reconstruction of tracks within these jets is crucial for many reconstruction and identification algorithms. In particular, flavour-tagging algorithms (Section 6.4.6), such as b-tagging, are dependent on the precise tracking information to reconstruct displaced vertices and measure track impact parameters. To ensure that these algorithms accurately model the behaviour observed in real collision data, it is important to quantify the fraction of tracks that are not reconstructed. This quantity, commonly referred to as the *fraction of lost tracks (*$\mathrm{f_{lost}}$*)*, is estimated using the energy loss (dE/dx) measurements of the charged clusters left by particles in the pixel detector. A track is considered lost when a single reconstructed track is associated with a cluster identified as originating from multiple charged particles. The quantity $\mathrm{f_{lost}}$ is measured separately in collision data and in MC simulations. A comparison between the results obtained in data and simulation allows an uncertainty to be assigned to the MC prediction. This uncertainty is propagated to the algorithms used to identify particles produced in these environments.

Figure 6.5 shows $\mathrm{f_{lost}}$ in the core of jets as a function of jet p_T, for Run-3 data collected in 2022 and the corresponding simulation[2]. It can be seen that $\mathrm{f_{lost}}$ increases with increasing jet p_T. This behaviour arises because, at higher jet p_T, the tracks become increasingly collimated, increasing the probability of track loss. The estimation of $\mathrm{f_{lost}}$ is done by fitting the dE/dx distributions in B-Layer pixel clusters using templates corresponding to single-particle and multiple-particle clusters [26]. The observed difference of approximately 25% between data and simulation likely originates from the limited accuracy with which the MC models the dE/dx distributions observed in data. This discrepancy is propagated as an uncertainty to algorithms affected by track loss, such as flavour-tagging algorithms.

[2] The author played a leading role in the development of the framework required to produce this result.

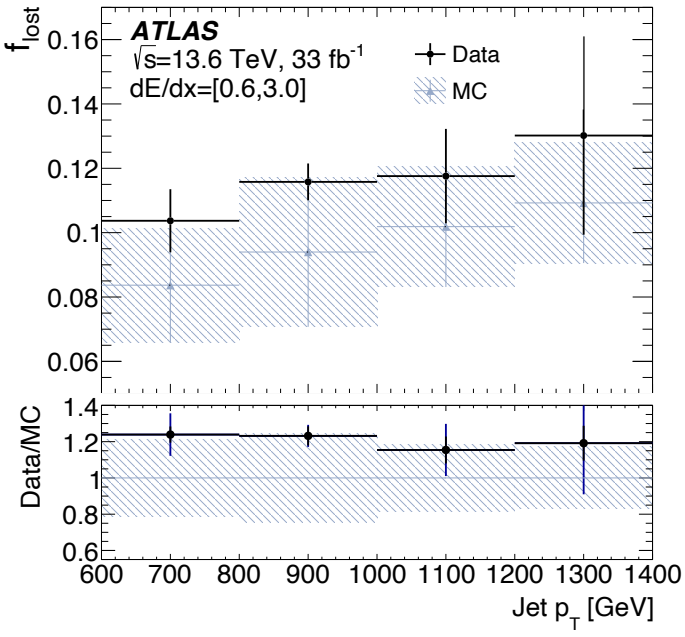


Figure 6.5: Results of f_{lost} as a function of jet p_T in Run-3 data collected in 2022 and the corresponding MC simulation. The measurement is based on dE/dx distributions in the range $0.6 < dE/dx < 3.0$ MeV g^{-1} cm^2. The error bands include both statistical and systematic uncertainties [26].

6.2 Electrons

The relatively low mass and electric charge of electrons make them highly interactive with the detector material in the electromagnetic calorimeter. As described in Section 4.2.2, the electrons induce electromagnetic showers in the electromagnetic calorimeter, where nearly all of the initial energy of the electron is deposited as clusters, with minimal residual deposits in the hadronic calorimeter. Electron reconstruction, thus proceeds by associating ID tracks with the corresponding energy deposits in the electromagnetic calorimeter [193]. Figure 6.6 shows a schematic illustration of the path of an electron through the ATLAS detector. The energy of the reconstructed electron candidates is subsequently calibrated to correct for detector effects and residual differences between data and simulation. Electron identification builds upon this reconstruction by exploiting discriminating variables that separate real electrons from backgrounds such as hadrons misidentified as electrons or electrons from photon conversions. Finally, electron isolation is applied to suppress non-prompt electrons originating from hadron decays or jets. Efficiency differences between data and simulation in the reconstruction, identification, and isolation steps are accounted for in simulated events through the application of scale factors.

Electron reconstruction In ATLAS, the energy deposits in the calorimeters are reconstructed as *topological clusters* or *topo-clusters* [195, 196]. The clustering algorithm is based on the cell

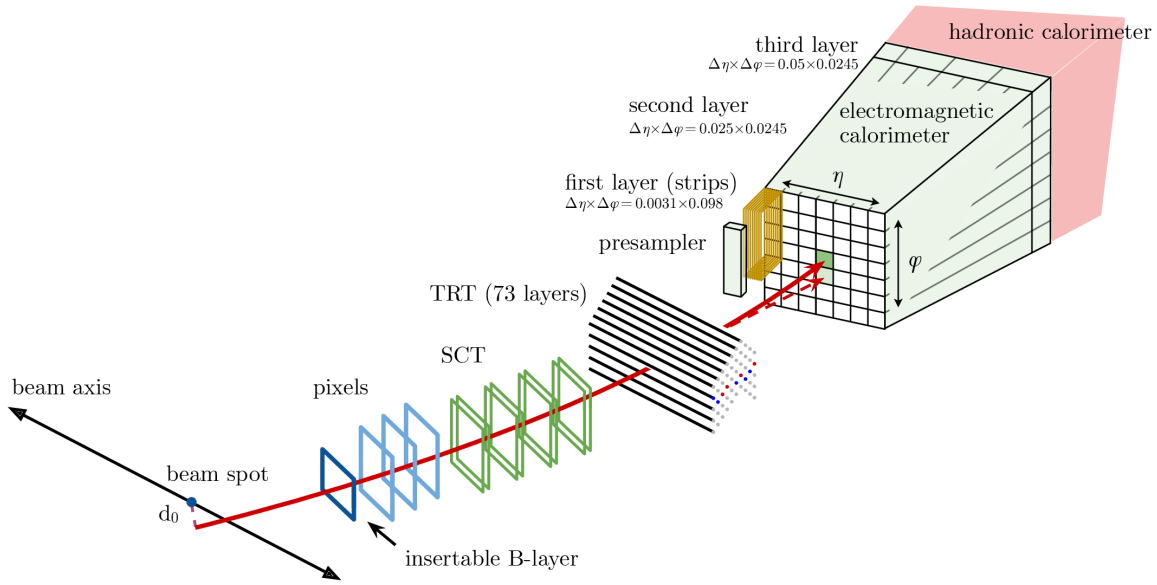


Figure 6.6: Illustration of the trajectory of an electron in the ATLAS detector. The red curve represents the path of the electron before it enters the electromagnetic calorimeter, while the dashed line indicates the trajectory of a photon radiated by the electron due to its interaction with the tracker material [194].

significance, defined as the ratio of the cell[3] energy (at the EM scale[4]) to the expected cell noise. Initially, *proto-clusters* are formed by merging neighbouring cells with cell significance above predefined thresholds [194]. Proto-clusters containing multiple local energy maxima above a given threshold are subsequently split into separate clusters. Only clusters with EM energy greater than 0.4 GeV are retained as topo-clusters, suppressing contributions from pile-up interactions.

Electron reconstruction then associates these topo-clusters with charged-particle tracks reconstructed in the ID, following the general track reconstruction and fitting procedure described in Section 6.1. While topo-clusters determine the electron energy, ID tracks provide information on the particle momentum and charge. To better account for bremsstrahlung energy losses in the detector material, tracks with at least four silicon hits and loosely matched to an EM cluster are refitted using a Gaussian-sum filter (GSF) [197], which is an extension of the Kalman filter that models the energy loss through a weighted sum of Gaussian components. The track-cluster matching is then done by imposing requirements on the separation in η and ϕ between the track and the cluster position.

Once the topo-clusters are matched to tracks, those with energies above 1 GeV and associated with tracks containing at least four silicon hits are selected. These serve as seeds for the construction of *superclusters*. Additional nearby clusters, referred to as *satellites*, are associated with the seed if they lie within predefined $\Delta\eta \times \Delta\phi$ windows around the seed barycentre. These satellites typically correspond to secondary EM showers produced by bremsstrahlung. The seed cluster and its satellites are combined to form a supercluster, to which calorimeter cells from the presampler and the first three LAr layers are assigned. In the transition region $1.4 < |\eta| < 1.6$, energy measured in the scintillator

[3] In the ATLAS calorimeters, a cell is the smallest readout unit of the detector.

[4] The EM scale is the basic calibration of the calorimeter signal that ensures the correct measurement of energy deposited by EM showers.

between the calorimeter cryostats is also included. An initial energy calibration and position correction are then applied to the supercluster. Tracks are matched to the resulting electron superclusters using criteria similar to those used for the initial track–cluster association.

Electron identification After reconstruction, it is essential to identify prompt electrons and separate them from objects that mimic their signature, including electrons from photon conversions in the detector material and non-prompt electrons from the decay of heavy-flavoured hadrons. ***Electron identification*** in ATLAS is performed using a *likelihood-based (LH) identification algorithm* [194, 198], which combines measurements from the tracking system, the calorimeter, and variables that link both subsystems. A discriminant is then constructed to distinguish prompt electrons (signal) from background objects. Based on the the discriminant value, three main ***working points (WPs)*** are defined: ***Loose***, ***Medium***, and ***Tight***, corresponding to progressively stricter selection criteria. The typical identification efficiencies for a prompt electron with $E_{\mathrm{T}} = 40\,\mathrm{GeV}$ are approximately 93%, 88%, and 80% for the Loose, Medium, and Tight WPs, respectively. The identification efficiencies are measured in data using the well-known processes $J/\psi \to e^+e^-$ and $Z \to e^+e^-$.

Electron isolation Even after applying a set of identification criteria, reconstructed electron candidates can originate from background processes, particularly from heavy-flavour decays or jets that produce electrons embedded within a dense environment of nearby particles. To suppress these backgrounds, an ***electron isolation*** requirement is applied [194, 198]. Isolation exploits the fact that prompt electrons from electroweak decays (such as from the decay of heavy resonances, including the Higgs, W, and Z bosons) are typically isolated, meaning they have little additional activity in their vicinity, whereas non-prompt or fake electrons from hadron decays or photon conversions are usually surrounded by significant hadronic energy or tracks. To quantify this effect, isolation variables are constructed to measure the amount of activity around the electron candidate. This is typically performed by summing the transverse energies of calorimeter clusters or the transverse momenta of tracks within a cone of radius $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$ around the direction of the electron candidate, excluding the candidate itself. Two classes of isolation variables are therefore defined: *calorimeter-based isolation* and *track-based isolation*, denoted by $E_{\mathrm{T}}^{\mathrm{iso}}$ and $p_{\mathrm{T}}^{\mathrm{iso}}$, respectively. Isolation WPs are defined by comparing these isolation variables to the transverse energy or momentum of the electron candidate, denoted by E_{T} and p_{T}, respectively. The WPs relevant for the analyses presented in this thesis are summarised in Table 6.1 and described in detail in Ref. [194]. For the track-based isolation, a variable-size cone, ΔR^{var}, is used. The cone size decreases with increasing electron p_{T} and is defined as $\min\left(\frac{10\,\mathrm{GeV}}{p_{\mathrm{T}}}, R_{\mathrm{max}}\right)$. In release 21, a maximum radius of $R_{\mathrm{max}} = 0.2$ is used, while the current recommendation for release 25 uses $R_{\mathrm{max}} = 0.3$. The latter configuration corresponds to the *Tight_VarRad* WP shown in Table 6.1. This variable-size cone approach accounts for the increasing collimation of decay products from highly boosted particles. The track-based isolation requirement is then imposed by requiring the sum of the transverse momenta of tracks within the cone to be smaller than 6% of the transverse momentum of the electron.

Electron calibration Reconstructed electron candidates are calibrated to correct the electron energy scale in data and the energy resolution in simulation. This ***calibration*** is derived in $Z \to e^+e^-$ events in data and simulation [200]. In addition, differences in reconstruction, identification, and isolation

Table 6.1: Definition of the electron isolations WPs relevant for this thesis. The first one is used in the threshold entanglement measurement, while the second one is used in the boosted analysis. Fixed-cone and variable-cone requirements are applied to the calorimeter- and track-based isolation variables, respectively [194, 198, 199].

WP	Calorimeter-based isolation	Track-based isolation
FCTight	$E_{\mathrm{T}}^{\mathrm{iso}}(\Delta R < 0.2)/E_{\mathrm{T}} < 0.06$	$p_{\mathrm{T}}^{\mathrm{iso}}(\Delta R^{\mathrm{var}} < 0.2)/p_{\mathrm{T}} < 0.06$
Tight_VarRad	$E_{\mathrm{T}}^{\mathrm{iso}}(\Delta R < 0.2)/E_{\mathrm{T}} < 0.06$	$p_{\mathrm{T}}^{\mathrm{iso}}(\Delta R^{\mathrm{var}} < 0.3)/p_{\mathrm{T}} < 0.06$

efficiencies between data and simulation must be accounted for. This is achieved by applying scale factors to simulated events, which are derived from comparisons of efficiencies measured in $Z \to e^+e^-$ and $J/\psi \to e^+e^-$ events in data and simulation [194].

6.3 Muons

Due to their relatively large mass and low interactive nature, muons traverse the calorimeters with minimal energy loss and are typically only identified in the ID and the MS. This property makes muons easily distinguishable from most other particle types and provides a clean experimental signature. Muon reconstruction in ATLAS combines information from the ID, the MS, and the calorimeters. After reconstruction, muon identification is performed to distinguish real muons from background objects, such as hadrons misidentified as muons or muons from in-flight decays of pions and kaons. To further suppress non-prompt muons originating from hadron decays or jets, an isolation requirement is applied [201–203].

Muon reconstruction Muon reconstruction in the ID follows the standard charged-particle tracking procedure, as described in Section 6.1. In the MS, reconstruction begins with the identification of short, straight track segments within individual stations (definition of stations in Section 4.2.4) using a Hough transform [204]. Segments from different stations are then combined into preliminary track candidates, assuming a parabolic trajectory consistent with the bending of muon trajectories in the magnetic field and a loose pointing constraint to the interaction point. Precision measurements in the bending plane are combined with second-coordinate information from the trigger detectors to form three-dimensional track candidates. A global χ^2 fit is then performed to determine the muon trajectory in the magnetic field, accounting for material interactions and possible detector misalignments. The resulting trajectory is then refined by removing outlier hits and including additional compatible ones. The fit is performed again with the updated hit information. Ambiguities are resolved by retaining the highest-quality tracks. Finally, tracks are refitted with a loose interaction point constraint, corrected for calorimeter energy loss, and extrapolated back to the beam line.

Reconstructed muons are grouped into different categories according to the subdetector that is primarily used in their reconstruction. The primary category consists of ***combined muons (CB)***, reconstructed by matching ID and MS tracks and performing a combined fit that accounts for calorimeter energy loss. ***Inside-out combined muons (IO)*** are built by extrapolating ID tracks to the MS and searching for loosely aligned MS hits, improving efficiency in regions with limited MS coverage or for low-p_{T} muons. When an MS track cannot be matched to an ID track, ***muon-spectrometer extrapolated muons (ME)*** are defined by extrapolating the MS track to the beam line, extending the

acceptance up to $|\eta| = 2.7$. ***Segment-tagged muons (ST)*** are identified when an ID track matches at least one MS segment, while ***calorimeter-tagged muons (CT)*** are obtained by associating ID tracks with calorimeter energy deposits consistent with a minimum-ionising particle [5]. Due to the large low-p_{T} background, CT muons require $p_{\mathrm{T}} > 5$ GeV, compared with 2 GeV for the other muon types.

Muon identification After reconstruction, muon identification algorithms are used to select high-quality muon candidates based on requirements on the number of hits in the different ID subdetectors and MS stations, the normalised χ^2 of the combined track fit, and variables testing the compatibility of measurements between the two detector systems. One such variable is the q/p compatibility, defined for CB and IO muons with an MS track as [201]:

$$q/p \text{ compatibility} = \frac{|q/p_{\mathrm{ID}} - q/p_{\mathrm{MS}}|}{\sqrt{\sigma^2(q/p_{\mathrm{ID}}) + \sigma^2(q/p_{\mathrm{MS}})}},$$

where q/p_{ID} and q/p_{MS} are the charge-to-momentum ratios measured in the ID and MS at the interaction point and $\sigma(q/p_{\mathrm{ID}})$ and $\sigma(q/p_{\mathrm{MS}})$ are the corresponding uncertainties.

A specific set of requirements on these variables for each muon type defines a selection WP, the three main ones being ***Loose***, ***Medium***, and ***Tight*** [201]. The *Loose* WP includes all muon types, with selection criteria optimised to maximise reconstruction efficiency, at the cost of slightly reduced purity. The *Medium* WP primarily selects CB and IO muons with good track quality and consistent ID–MS measurements, providing a balanced choice between efficiency and background rejection across the full detector acceptance. The *Tight* WP applies the strictest selection, accepting only CB and IO muons with high-quality tracks, optimised q/p criteria, and a combined track fit χ^2 cut, achieving maximum background rejection while maintaining high efficiency for prompt muons.

Muon isolation Muon isolation is defined in analogy to that used for electrons, using both calorimeter-based and track-based isolation variables. The isolation cone for muons has a maximum radius of $\Delta R^{\mathrm{var}} = 0.3$. In analogy with electrons (Table 6.1), only the track-based variable is used for muons in the threshold entanglement analysis. This requirement defines the *FCTightTrackOnly* isolation WP, summarised in Table 6.2. For the boosted regime, both calorimeter- and track-based isolation variables are employed, defining the *Tight_VarRad* WP, with slightly modified requirements compared to the electron definitions.

Table 6.2: Definition of the muon isolation WPs used in this thesis. The first one is used in the threshold entanglement measurement, while the second one is used in the boosted analysis.

WP	Calorimeter-based isolation	Track-based isolation
FCTightTrackOnly	-	$p_{\mathrm{T}}^{\mathrm{iso}}(\Delta R^{\mathrm{var}} < 0.3)/p_{\mathrm{T}} < 0.06$
Tight_VarRad	$E_{\mathrm{T}}^{\mathrm{iso}}(\Delta R < 0.2)/E_{\mathrm{T}} < 0.15$	$p_{\mathrm{T}}^{\mathrm{iso}}(\Delta R^{\mathrm{var}} < 0.3)/p_{\mathrm{T}} < 0.04$

[5] A minimum-ionising particle (MIP) is a charged particle traversing matter with an energy loss near the minimum of the Bethe–Bloch curve [31], depositing a small and approximately constant amount of energy per unit length. Muons in the GeV range behave as MIPs, losing negligible energy in the calorimeter and thus reaching the MS in the ATLAS detector.

Muon calibration Similar to electrons, muon candidates are calibrated to correct the simulation for differences in the momentum scale and resolution relative to data. This calibration is derived from resonance decays, specifically $Z \to \mu^+\mu^-$ and $J/\psi \to \mu^+\mu^-$ events. Muon reconstruction, identification and isolation efficiencies in MC are also corrected by applying scale factors derived using efficiencies measured in data for $Z \to \mu^+\mu^-$ and $J/\psi \to \mu^+\mu^-$ events [203].

6.4 Jets

When two protons collide, it is in fact their constituent partons that participate in the hard scattering. As discussed in Section 2.1.2, the outgoing partons, predominantly quarks and gluons, cannot exist freely due to colour confinement and therefore hadronise, forming colour-neutral particles known as hadrons. The resulting hadrons are produced in collimated sprays along the direction of the parton, forming bundles of particles called ***jets***. Jets serve as experimental proxies for the original quarks or gluons, with their reconstructed four-momentum ideally representing, as precisely as possible, the kinematics of the parent partons.

Jets are reconstructed using ***jet clustering algorithms***, which use as inputs different types of physics objects. Depending on the choice of input objects and the algorithm used, various jet collections can be defined within a single collision event. The measured jet energy, however, can be affected by multiple factors such as detector response, energy losses, and pile-up effects. ***Jet calibration*** refers to the process of correcting the reconstructed quantities so that the measured jet energy, momentum, and direction accurately represent those of the original parton. In ATLAS, this is achieved through a sequence of dedicated corrections that bring the detector response to uniform and well-defined energy scales and directions, ensuring consistency between data and simulation.

6.4.1 Input objects for reconstruction

The reconstruction of jets begins with the formation of input objects, which in this case represent the detector-level observables used by the jet clustering algorithms. Depending on the reconstruction scheme, these inputs can be derived from calorimeter information, tracking detectors, or a combination of both. During Run 1, most ATLAS analyses used the so-called *EMTopo jets*, which are reconstructed using topo-clusters (Section 6.2) as input objects [195], corrected by an energy-scale factor [205]. Alternatively, jets can also be reconstructed from tracks measured in the ID. These track jets provide precise measurements of direction and momentum, which are particularly valuable for mitigating pile-up effects and improving jet energy resolution. To exploit the complementary strengths of the calorimeter and tracking subsystems, an ensemble combining simultaneously tracks from the hard interaction and topo-clusters can be used. This unified set of input objects are known as *particle flow objects (PFlow objects)* [206]. Jets reconstructed from PFlow objects are referred to as ***PFLow jets***.

The PFlow algorithm begins by matching ID tracks to topo-clusters in the calorimeter, ensuring that the energy and momentum from the same particle are not double-counted. A cell-based energy subtraction procedure is then performed where the expected calorimeter energy deposit corresponding to each matched track is computed from its momentum and impact position, and this energy is subtracted cell-by-cell from the associated topo-cluster. Tracks used in this algorithm are selected using stringent quality criteria, typically $|\eta| < 2.5$ and $0.5 < p_{\mathrm{T}} < 40\,\mathrm{GeV}$, while those associated with identified muons or electrons are excluded to avoid overlaps. In some cases, a single charged

particle may deposit energy across multiple topo-clusters. The algorithm evaluates the probability of such energy sharing and, when necessary, includes additional clusters to recover the full shower energy. Once the subtraction step is complete, the residual energy in the system is compared to the expected fluctuations from a single-particle shower. If consistent, the cluster remnants are removed. The resulting PFlow objects are then classified based on their composition: topo-clusters associated with tracks are identified as charged hadrons, while topo-clusters without matching tracks are assumed to originate from neutral hadrons.

Using PFlow objects as inputs for jet reconstruction provides several advantages over calorimeter-only approaches. By combining the precise tracking information from the ID with calorimeter measurements, the overall jet energy resolution is improved, particularly at low energies where tracking dominates. The inclusion of tracks with low-p_{T} that do not reach the calorimeter or fall below the calorimeter thresholds ensures a more complete reconstruction of jet constituents. Furthermore, pile-up effects are significantly reduced, as tracks can be associated with reconstructed vertices and those originating from pile-up interactions can be excluded.

6.4.2 Jet clustering algorithms

To reconstruct jets, various algorithms have been developed over the years [207–209] that use inputs, discussed in Section 6.4.1, to form jet candidates. Depending on the analysis strategy and the targeted final state, the jet definition and consequently the choice of clustering algorithm may vary.

An effective algorithm should satisfy several key requirements. It must be *infrared safe*, meaning that the reconstruction is unaffected by the addition of soft particles with negligible energy, and *collinear safe*, ensuring that the splitting of an input constituent into two collinear ones does not alter the reconstructed jet. Furthermore, the algorithm should minimise sensitivity to detector effects such as noise and resolution, while remaining computationally efficient.

The most common jet clustering algorithms employed at the LHC are based on sequential recombination, in which input objects or proto-jets are iteratively combined to form jets. The combination is determined by two distance metrics:

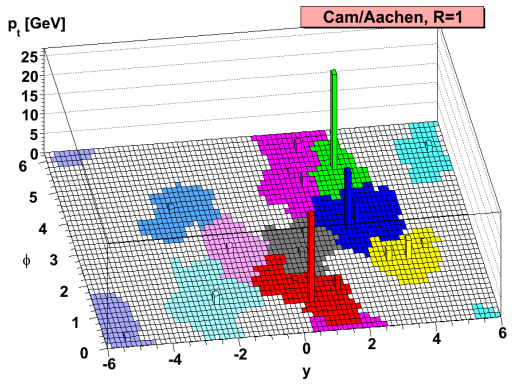
$$d_{ij} = \min(p_{T,i}^{2n}, p_{T,j}^{2n}) \frac{\Delta R_{ij}^2}{R^2},$$

$$d_{iB} = p_{T,i}^{2n},$$

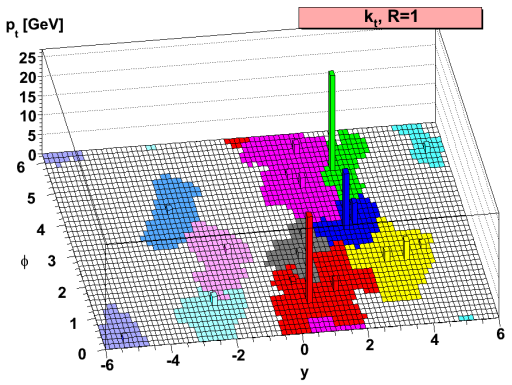
where $\Delta R_{ij} = \sqrt{(y_i - y_j)^2 + (\phi_i - \phi_j)^2}$ is the angular distance between two input objects, with y_i (y_j) and ϕ_i (ϕ_j) denoting the rapidity and azimuthal angle of particle i (j). The parameters n and R characterise the specific jet algorithm: n defines the *jet clustering strategy*, determining how proto-jets are merged, while R is a *radius parameter* representing the characteristic size of the jet. In ATLAS, two common choices are $R = 0.4$ and $R = 1.0$. The quantity $p_{T,i}$ denotes the transverse momentum of the i-th proto-jet. The quantity d_{ij} represents the distance between the i-th and j-th inputs, while d_{iB} defines the distance between the i-th input and the beam. For each input i, both d_{ij} and d_{iB} are computed relative to all other constituents in the event. If $d_{iB} > d_{ij}$, the two input objects i and j are combined, otherwise i is identified as a final jet and removed from the list of inputs. This procedure is repeated iteratively until all inputs in the event are clustered.

Different choices of n define distinct algorithms. For $n = 0$, one obtains the ***Cambridge-Aachen***

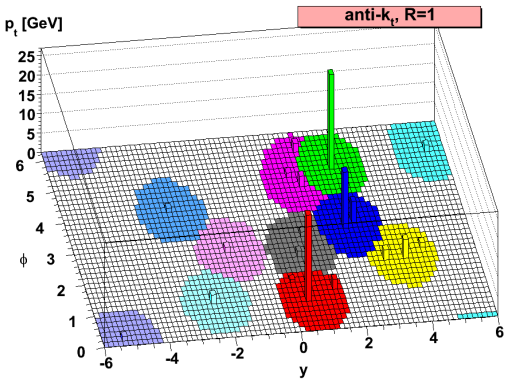
algorithm [207], where the recombination depends solely on the angular separation ΔR_{ij}, independent of transverse momentum. Objects with the smallest angular separation are merged first. For $n = 1$, the ***k_T algorithm*** is recovered [208]. Here, the clustering depends on both angular separation and transverse momentum, such that soft and nearby objects are merged first. This produces irregularly shaped jets that reflect the soft radiation pattern of the event, a behaviour generalised for any $n > 0$. Finally, for $n = -1$ the ***anti-k_T algorithm*** is obtained [209], which is the most widely used jet algorithm within ATLAS. In this case, hard and well-separated inputs are clustered first, establishing the jet axes early in the iteration. Softer nearby constituents are then added around these cores, producing jets with approximately conical shapes of radius R. Figure 6.7 illustrates the typical jet morphologies resulting from different jet clustering algorithms for the same event in the $y - \phi$ plane. Here, the z-axis denotes the p_T of the reconstructed jet. Unlike the Cambridge/Aachen and k_T algorithms, jets reconstructed with the anti-k_T algorithm exhibit an approximately circular shape in the $y - \phi$ plane.



(a) Cambridge/Aachen algorithm



(b) k_T algorithm



(c) Anti-k_T algorithm

Figure 6.7: Illustrative examples of reconstructed jet shapes obtained with different jet clustering algorithms in the $y - \phi$ plane. The z-axis denotes the p_T of the reconstructed jet. For simplicity, in all cases the radius parameter is set to 1.0. Different colours correspond to different jets. [209].

Unless stated otherwise, the analyses presented in this thesis use the anti-k_T algorithm for jet

reconstruction. This algorithm is infrared-safe, since clustering begins from hard objects before incorporating soft constituents, and collinear safe, as inputs with $R_{ij} < R$ are clustered together into a single jet.

6.4.3 Small-R jets

The default choice for jets in ATLAS analyses are those reconstructed using the anti-k_T algorithm with a radius parameter of $R = 0.4$, built from PFlow objects as inputs. In this thesis, these jets are referred to as ***small-R jets***. To reduce the impact of pile-up and to correct the reconstructed jet four-momentum, several correction and calibration procedures are applied, as detailed in Section 6.4.5. In addition, further mitigation of pile-up effects is achieved using the ***jet-vertex-tagging (JVT)*** technique, based on a multi-dimensional likelihood [210]. The JVT algorithm is used to identify jets originating from partons not associated with the hard-scattering primary vertex or with a significant pile-up contribution. It relies on the *corrected jet vertex fraction (corrJVT)* variable to determine whether a jet originates from a hard-scattering vertex or from pile-up. The corrJVF quantifies the relative contributions of tracks associated with the primary hard-scattering vertex and those from pile-up vertices within the same bunch crossing. Its values range from 0 to 1, with higher values corresponding primarily to jets originating from the hard-scattering interaction. In addition, the variable R_{p_T} is defined to measure the overall fraction of charged-particle transverse momentum within the jet that originates from the primary vertex. This variable typically peaks at zero for pile-up jets with no associated primary-vertex tracks and is a steeply falling distribution for pile-up jets. Both corrJVF and R_{p_T} are combined into a single discriminant, referred to as the JVT, constructed as a two-dimensional likelihood derived from simulated dijet events and implemented using a k-nearest neighbour (KNN) algorithm [211]. A selection requirement is then placed on the value of the JVT discriminant to select jets originating from the main hard-scattering vertex. Different threshold choices correspond to different efficiencies for retaining jets in an event, and these are referred to as working points.

6.4.4 Reclustered jets

When heavy particles such as top quarks, W bosons, or Higgs bosons decay, the angular separation between the decay products depends on the transverse momentum of the parent particle. For a two-body decay $X \to a + b$, this separation is approximately given by [212]:

$$\Delta R(a, b) \sim \frac{2m^X}{p_T^X},$$

where $\Delta R(a, b)$ denotes the angular separation in the $y - \phi$ plane between the two decay products a and b of a heavy particle X with mass m^X and transverse momentum p_T^X. The higher the transverse momentum of the heavy particle, the smaller the angular separation. In such cases, the decay products of the heavy particle are too close to be reconstructed as distinct small-R jets, using the procedure described in Section 6.4.3. Instead, they are typically reconstructed within a single large-radius (large-R) jet.

In the context of this thesis, this effect is particularly relevant for top quarks produced in the high $m_{t\bar{t}}$ regime. In this regime, they acquire large Lorentz boosts and their hadronic decay products become highly collimated. Therefore, in the boosted entanglement analysis, the hadronically decaying

top quark is reconstructed as a single large-R jet.

Various approaches can be used to reconstruct such jets, depending on the choice of inputs and clustering algorithm. In principle, a large-R jet could be reconstructed directly from PFLow objects with a radius parameter of $R = 1.0$, in analogy to small-R jet reconstruction. However, this would require an independent jet clustering approach with a dedicated calibration and optimisation of jet substructure observables. This approach is therefore not studied in the scope of this thesis. Instead, a *reclustering technique* is employed, in which calibrated $R = 0.4$ anti-k_T PFLow jets are used as inputs to construct a large-R jet with a fixed radius of $R = 1.0$. This jet, referred to as a ***reclustered (RC) jet***, is used to represent the hadronically decaying top quark candidate in the boosted regime. An illustration of such a scenario is shown in Figure 6.8. Reclustering offers the advantage of relying on input jets that are already fully calibrated and therefore do not require an additional large-R jet calibration [213]. This provides greater flexibility in the choice of parameters in the clustering algorithm.

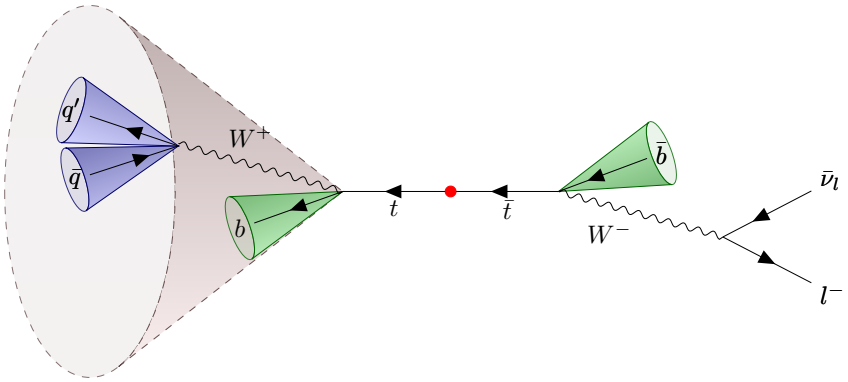


Figure 6.8: Diagram illustrating the topology of a $t\bar{t}$ event in the lepton+jets decay channel containing a boosted top quark. The hadronically decaying top quark is reconstructed as a RC large-R jet (shown in pink), which contains two small-R jets (blue) originating from the hadronic decay of the W boson and a small-R jet identified as a b-jet via b-tagging (green).

6.4.5 Jet calibration

Reconstructed jets are subject to various detector and reconstruction-related effects that can distort their measured energy and direction. These effects arise from the imperfect calorimeter response to different particle types, energy losses in inactive material, and additional energy deposits from pile-up interactions. Consequently, the reconstructed jet energy does not directly correspond to the true energy of the particle-level jet originating from the initial parton. To account for these effects, jet energy scale calibration is performed. The main purpose of the calibration is to restore the ***jet energy scale (JES)*** and direction to those of the corresponding particle-level jet, ensuring consistency between data and simulation. It is also important to optimise the overall ***jet energy resolution (JER)*** to allow for a precise reconstruction of the jet kinematic properties. A number of uncertainties are associated with the different calibration methods. Details of these uncertainties are available in Section 7.9.

Jet energy scale calibration

The calibration process typically comprises several stages that correct the four-momentum by scaling the jet p_T, energy, and mass [214]. All these calibrations are applied to both data and MC simulations, unless otherwise stated:

Pile-up corrections The first step in the jet calibration chain is the correction for pile-up, which removes excess energy arising from additional pp interactions from in-time or out-of-time bunch crossings. These corrections consist of two components. The primary correction is based on the *jet area*, A and the *event transverse momentum density*, ρ, in the $y - \phi$ plane. The jet area quantifies how susceptible a jet is to pile-up and in this case it is determined from the number of infinitely soft particles (*ghosts*) clustered into the jet [215]. The quantity ρ is calculated as the median transverse momentum density of jets reconstructed with the k_T algorithm ($R = 0.4$) in the central region ($|\eta| < 2$). The k_T algorithm is chosen for its tendency to cluster a uniform soft background, while the use of the median value reduces sensitivity to hard-scatter jets. The pile-up contribution to the jet is calculated as the product ρA and is subtracted from the reconstructed jet p_T.

Since the ρ calculation is derived from the central region of the calorimeter, a residual pile-up correction is applied to account for any remaining dependence of the reconstructed jet p_T on pile-up activity. This dependence is parametrised in terms of the number of reconstructed primary vertices, N_{PV}, and the average number of interactions per bunch crossing, μ. The corrected jet p_T, after both the p_T-density-based and residual corrections, is given by [214]:

$$p_{\text{T}}^{\text{corr}} = p_{\text{T}}^{\text{reco}} - \rho \times A - \alpha \times (N_{\text{PV}} - 1) - \beta \times \mu,$$

where $p_{\text{T}}^{\text{reco}}$ and $p_{\text{T}}^{\text{corr}}$ denote the reconstructed jet p_T before and after pile-up correction, respectively. The parameters α and β are obtained from comparisons between reconstructed and truth-level jets.

Jet energy scale and η calibration The JES defines the relationship between the measured energy of a reconstructed jet and the energy of the corresponding particle-level jet. The ***JES calibration*** is applied to restore the reconstructed jet energy to that of the particle-level jet. This correction is derived by measuring the average jet energy response,

$$R = \frac{E^{\text{reco}}}{E^{\text{true}}},$$

where E^{reco} is the reconstructed jet energy and E^{true} is the jet energy at particle level. The calibration is parameterised as a function of E^{true}, and the detector pseudorapidity, η_{det}, which denotes the pseudorapidity of the jet direction measured with respect to the geometrical centre of the detector rather than the event vertex. The response is evaluated as the mean of a Gaussian fit to the core of the $E^{\text{reco}}/E^{\text{true}}$ distribution, in bins of E^{true} and η_{det}. The average response is parametrised as a function of E^{reco} using a numerical inversion procedure [216] and the jet energy calibration factor is then defined as the inverse of the average energy response.

In addition, a dedicated ***η calibration*** is applied to correct for small biases in the reconstructed jet direction by aligning the reconstructed pseudorapidity, η_{reco}, with the true value, η_{truth}. This correction is parameterised as a function of E^{true} and η_{det}, and it corrects the jet p_{T} and η components of the four-momentum without modifying the full four-vector.

Global sequential calibration Even after applying the above calibrations, the jet response still exhibits jet-to-jet fluctuations due to differences in the composition of particles, shower development, and particularly the parton flavour (quark or gluon) of the initiating parton. Quark-initiated jets tend to have harder particles and deeper calorimeter penetration, while gluon-initiated jets are broader with softer constituents, leading to a lower detector response. To reduce these effects and improve jet resolution without changing the average jet energy response, a ***global sequential calibration (GSC)*** is applied [214]. This procedure applies a set of multiplicative corrections based on global jet observables such as calorimeter shower shapes, tracking information, and muon activity. Independent corrections are derived for each of these observables using a numerical inversion method, in bins of jet p_T and η_{det}.

In-situ jet calibrations The final stage of jet energy calibration is used to correct residual differences in jet response between data and simulation. Unlike previous calibrations, this correction is applied only to data. Jet responses are measured independently in data and MC simulation and their ratio is used as an additional correction to data.

Three in-situ calibrations are used. The η calibration aligns the response of forward jets ($0.8 \leq |\eta_{\text{det}}| < 4.5$) to that of central jets ($|\eta_{\text{det}}| < 0.8$). The second calibration uses Z+jet and γ+jet analyses to measure the balance between the hadronic recoil and a well-calibrated Z boson or photon. Finally, the multijet balance technique calibrates a high-p_T jet using a system of low-p_T jets [214].

Jet energy resolution

Once all the above-mentioned corrections are applied, the JER is evaluated in bins of p_T and η. The JER quantifies the relative spread of the reconstructed jet momentum with respect to its true value and can be parameterised as [214]:

$$\frac{\sigma(p_T)}{p_T} = \frac{N}{p_T} \oplus \frac{S}{\sqrt{p_T}} \oplus C.$$

It has three independent contributions. The *noise term (N)* arises from detector electronics and pile-up effects, the *stochastic term (S)* reflects statistical fluctuations in the energy deposited in the calorimeter, and the *constant term (C)* accounts for residual effects linked to imperfect calibration. The JER is measured using an asymmetry-based dijet balance method, while the contribution from the noise term is determined separately using the random cones method [214]. A combined JER measurement is obtained by fitting the dijet balance results with a constraint on the noise term from the random-cone measurement. The procedure is repeated in η bins using the functional form above. Further details can be found in Ref. [214].

6.4.6 Flavour tagging

Once jets are reconstructed using the relevant jet algorithms, some analyses require additional information about their origin, in particular the flavour of the quark from which the jet may have originated. ***Tagging*** refers to the identification of jets originating from quarks of a given flavour or from gluons. The analyses discussed in this thesis involve the decay of top quarks into b-quarks and W bosons, where the b-quarks subsequently hadronise to form jets. Among all jets reconstructed in a given collision event, it is therefore important to identify those originating from b-quarks. For the analyses presented in this thesis, it is equally important to enhance the purity of events with jets that originate

from c-quarks. Various algorithms to identify such jets, referred to as ***flavour-tagging algorithms***, have been developed for this purpose. For Run-2, one of the most commonly used algorithms to identify b-jets is the deep-neural network based ***DL1r b-tagging algorithm***. More efficient algorithms based on more modern machine learning tools are becoming available for Run-3 analyses. One such algorithm is the transformer-based ***GN2 flavour tagging algorithm***. This algorithm is only compatible with newer releases and so in this thesis has only been used in the release 25-based boosted analysis.

DL1r algorithm

Unlike hadrons containing lighter quarks, b-hadrons have relatively long lifetimes of about 1.5×10^{-12} s [31], primarily due to the structure of the CKM matrix (Section 2.1.3). As a result, they typically travel a distance of approximately 450 μm, before decaying into lighter hadrons. As a result, b-jets typically contain a secondary vertex, and tracks exhibiting large impact parameters, as illustrated in Figure 6.9. The DL1r b-tagging algorithm exploits this information to discriminate b-jets from jets originating from c-quarks or light-flavour quarks. The algorithm is based on fully connected feed-forward neural networks [217]. It uses as inputs the outputs of several intermediate algorithms, which are designed to identify secondary vertices and large track impact parameters [218]. In addition, the outputs of the RNNIP algorithm are included as inputs. RNNIP is a recurrent neural network that processes a variable number of tracks associated with the jet. It uses track-level information such as the fraction of the jet p_{T} carried by each track, the angular distance between the track and the jet axis, and quantities related to the impact parameter of the track with respect to the primary vertex. It the outputs jet flavour probabilities that encode the likelihood that a jet originates from a b-, c-, or light-flavour quark.

All the inputs to the DL1r algorithm are then combined to estimate the probabilities p_b, p_c, and p_{light} for a jet to originate from a b-, c-, or light-flavour quark, respectively. These probabilities are combined into a single DL1r b-jet discriminant, defined as [218]:

$$\mathrm{DL1r} = \log\left(\frac{p_b}{f_c p_c + (1 - f_c) p_{\mathrm{light}}}\right), \tag{6.1}$$

where the parameter f_c represents the assumed fraction of c-jets in the background and can be tuned to optimise rejection of both c- and light-flavour jets. Figure 6.10 shows the distribution of the DL1r b-jet discriminant output for b-, c-, and light-flavour jets in $t\bar{t}$ simulated events calculated in Ref. [218]. It can be seen that the b-jets are predominantly concentrated at higher discriminant values, while light-favour jets populate the lower values. There also exists a region where c-jets tend to dominate. Based on this behaviour, the performance of the algorithm for b-tagging is characterised by the so-called ***b-tagging efficiency***, defined as the efficiency for correctly identifying a jet originating from a b-quark.

Different selection thresholds, or cut values, applied to the DL1r discriminant correspond to different average b-tagging efficiencies and background rejection rates. In ATLAS, standard WPs are defined for b-tagging efficiencies of 60%, 70%, 77%, and 85%. These WPs correspond to progressively lower thresholds on the DL1r discriminant. The tightest WP (60%) provides a high purity of b-jets but a lower tagging efficiency, whereas the loosest WP (85%) yields a higher efficiency at the cost of reduced purity.

Since the DL1r b-jet discriminant is a continuous variable, it can be divided into discrete intervals

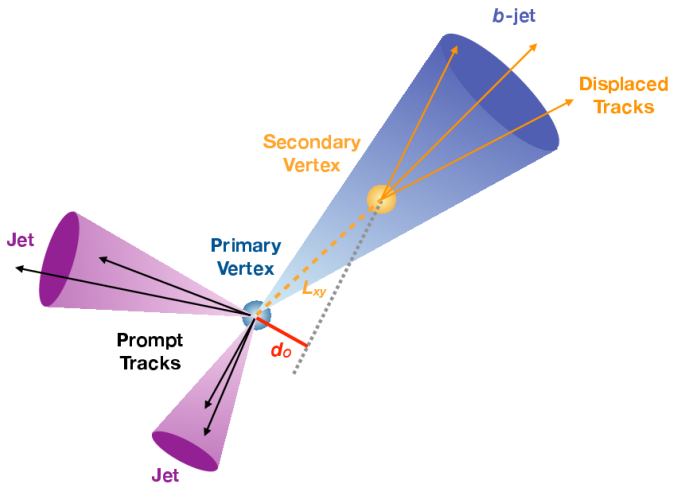


Figure 6.9: Illustration of the secondary vertex, decay length (L_{xy}), and the transverse impact parameters d_0 for a b-hadron decay within a jet emanating from the primary interaction vertex [219].

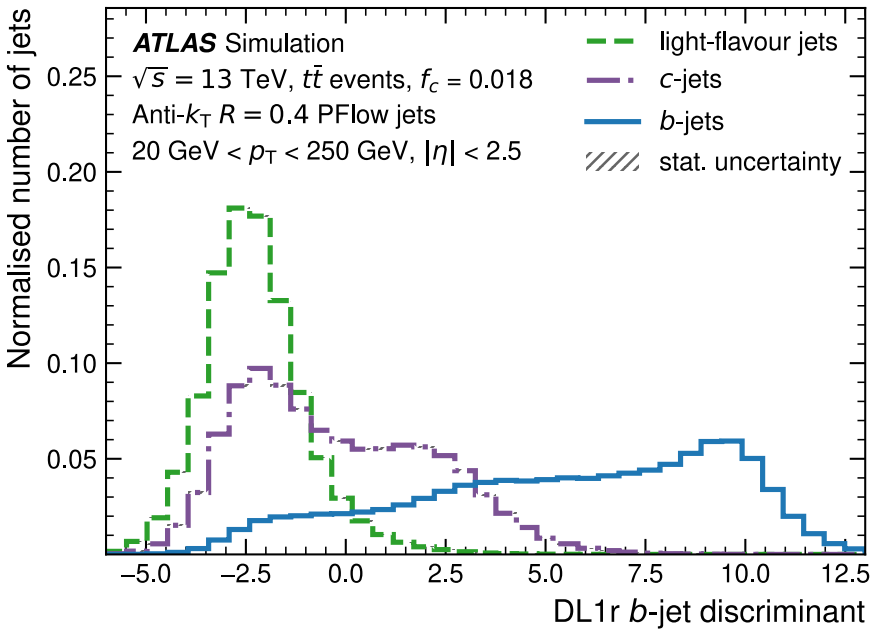


Figure 6.10: Distribution of the DL1r b-jet discriminant output for b-jets (solid), c-jets (dashed-dotted), and light-flavour jets (dashed). The curves represent the normalised probability distributions of the DL1r b-jet discriminant output for each jet flavour, i.e. the probability of a jet of a given flavour to obtain a particular b-tagging discriminant value. The distribution peaks at higher values for b-jets, while c- and light-flavour jets are concentrated at lower values, with an intermediate enhancement region visible for c-jets [218].

defined based on the cut values adopted for each WP. These intervals correspond to b-jet tagging efficiencies of [100, 85]%, [85, 77]%, [77, 70]%, [70, 60]%, and [60, 0]%. Most of the light-flavour jets fall within the [100, 85]% bin, while the [60, 0]% bin is enriched in b-jets [220]. Each interval is assigned an integer score, forming the so-called pseudo-continuous b-tagging (PCBT) scores. This variable provides a simplified, binned representation of the discriminant output that retains approximate flavour separation power. The PCBT score is used as an input to the tool developed for the reconstruction of the $t\bar{t}$ system (Section 7.4), and also to define the charm-enriched signal region described in Section 7.6.

Analogous to Equation 6.1, a DL1r c-jet discriminant can also be defined [218]. In the entanglement analysis, the identification of c-jets is particularly important, as it allows for the indirect identification of the down-type jet used as the spin analyser. However, there are currently no calibrated DL1r ***c-tagging*** WPs available in release 21. For the threshold entanglement analysis presented in this thesis, a dedicated c-tagging strategy is therefore not employed. Instead, a charm-enriched signal region is constructed using the available PCBT scores themselves, as described in Section 7.6. This approach presents the most practical option with the current tagging capabilities, without introducing an additional dedicated c-tagging calibration.

Mismodelling of the DL1r discriminant distribution in MC simulations leads to discrepancies in the observed b-tagging rates between data and simulation. To account for these differences, scale factors derived from data are applied to MC simulated events, ensuring that the b-tagging efficiency in MC matches that measured in data [220]. This b-tagging algorithm is applied exclusively to small-R jets, and typically calibrated for anti-k_T $R = 0.4$ PFLow jets.

GN2 algorithm

With advances in machine learning, more robust flavour-tagging algorithms are being developed which achieve higher background rejection at the same signal efficiency. One such algorithm is GN2, which employs a transformer-based architecture [221]. In contrast to DL1r, GN2 does not rely on intermediate taggers. Instead it is based on the earlier GN1 algorithm [222], replacing the graph neural network with a transformer architecture together with several other additional optimisations. This model directly accepts information about jets and their associated tracks provided by the standard event reconstruction. A key characteristic of the GN2 model lies in the use of physics-inspired constraints that support the primary task of jet flavour classification while also improving the interpretability of the model. This is achieved by incorporating two additional training objectives: predicting the origin of tracks associated with the jet and identifying tracks that originate from a common vertex. Since these tasks are not strictly required for jet flavour classification, they are referred to as auxiliary training objectives. Furthermore, the GN2 is trained to identify an additional class of jets that originate from hadronic τ-lepton decays. The b-tagging discriminant, D_b, is therefore defined as:

$$D_b = \log\left(\frac{p_b}{f_c p_c + f_\tau p_\tau + (1 - f_c - f_\tau) p_u}\right),$$

where p_b, p_c, p_τ, and p_u represent the predicted probabilities for a jet to originate from a b-quark, c-quark, τ-lepton, or light-flavour parton, respectively. A jet is considered b-tagged if D_b exceeds a chosen threshold. A selection on D_b defines a WP with a certain b-tagging efficiency. Based on this, there are five WPs corresponding to b-tagging efficiences of 65%, 70%, 77%, 85%, and 90%. The

mis-tagging rate for c-, τ-, and light-flavour jets is defined as the fraction of jets of a given flavour that are incorrectly identified as b-jets. The reciprocal of this mistagging rate is referred to as the rejection factor for that jet flavour.

The performance of a b-tagging algorithm is evaluated based on its ability to reject c-, τ-, and light-flavour jets while maintaining a desired b-jet tagging efficiency. Figure 6.11 illustrates the performance of the tagger in terms of the c-jet, τ-jet, and light-jet rejection as a function of the b-jet tagging efficiency for simulated $t\bar{t}$ events. The GN2 algorithm provides significantly improved background rejection compared with DL1d across the full range of b-jet tagging efficiencies. Here, DL1d is the successor to the DL1r b-tagging algorithm. The main difference in architecture is that DL1d processes track parameters using the DIPS algorithm, which is based on the DeepSets framework [223], instead of a recurrent neural network. Owing to its improved efficiency and background rejection, GN2 algorithms currently represent the most powerful flavour-tagging algorithms available for the latest ATLAS analyses.

Furthermore, unlike DL1r, the GN2 framework now provides calibrated c-tagging WPs, enabling direct c-tagging. Analogously to the b-tagging discriminant, a dedicated c-tagging discriminant is defined, where only those jets failing the 77% b-tagging WP are considered. This selection ensures that there is no overlap between c-tagged jets and jets passing the tighter 77% b-tagging requirement, which are therefore less likely to originate from c-quarks. Based on this c-tagging discriminant, calibrated c-tagging WPs with efficiencies of 10%, 30%, and 50%, have been defined. This tagging strategy is particularly useful for the indirect identification of the down-type jet used as the spin analyser in entanglement measurements. This approach is planned to be implemented in the boosted entanglement measurement presented in this thesis, as discussed in Chapter 8.

6.5 Missing transverse momentum

Some particles produced in the pp collisions may escape detection in the ATLAS detector subsystems. Their presence can be quantified by a variable known as the ***missing transverse momentum***, denoted by $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$, with its magnitude referred to as $E_{\mathrm{T}}^{\mathrm{miss}}$. It serves as a proxy for undetected particles, such as neutrinos or possible non-interacting particles prediced by new physics models. The vector $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ is defined as the negative vectorial sum of the transverse momenta of all physics objects in an event. This definition is based on momentum conservation. In pp collisions, the initial transverse momentum of the system is approximately zero. Therefore, the vector sum of the transverse momenta of all final-state particles should also vanish. Any observed imbalance is attributed to particles that do not interact with the detector.

The reconstruction of $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ in ATLAS involves two components [224]. The first one is from fully reconstructed and calibrated physics objects called the hard-event signals associated with the primary vertex of interaction. The second contribution comes from soft-event signals characterised by reconstructed charged-particles tracks not associated with the hard-scatter vertex nor with any selected analysis object. They account for low-p_{T} activity in the event. The Cartesian components of $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ are given by:

$$E_{x(y)}^{\mathrm{miss}} = -\sum_{i \in \text{hard objects}} p_{x(y),i} - \sum_{j \in \text{soft signals}} p_{x(y),j}.$$

From these, the magnitude and azimuthal direction of the missing transverse momentum are

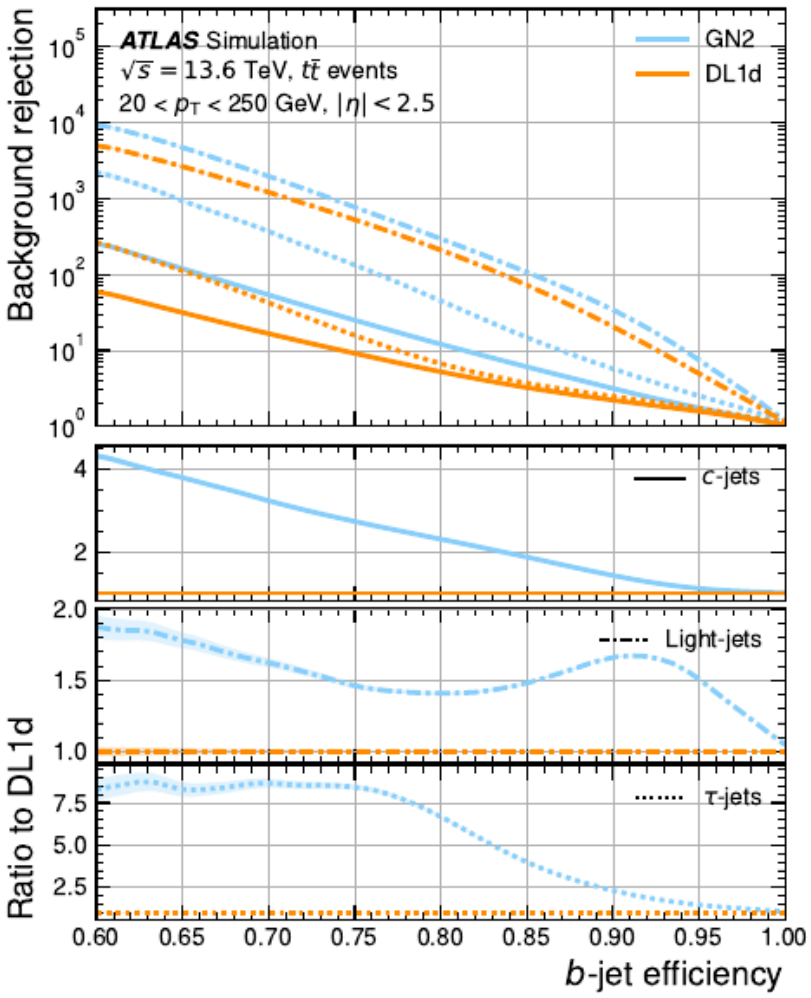


Figure 6.11: Rejection rates for c-jet (solid), light-jet (dash–dotted), and τ-jet (dashed) as a function of the b-jet tagging efficiency for jets in simulated $t\bar{t}$ events with $20 < p_T < 250$ GeV. Results are shown for the GN2 (light blue) and DL1d (dark orange) taggers. The lower panels display the relative performance of GN2 with respect to DL1d. Figure taken from Ref. [221].

obtained as:

$$\vec{E}_{\mathrm{T}}^{\mathrm{miss}} = (E_x^{\mathrm{miss}}, E_y^{\mathrm{miss}}),$$
$$E_{\mathrm{T}}^{\mathrm{miss}} = |\vec{E}_{\mathrm{T}}^{\mathrm{miss}}| = \sqrt{(E_x^{\mathrm{miss}})^2 + (E_y^{\mathrm{miss}})^2},$$
$$\phi^{\mathrm{miss}} = \tan^{-1}\left(\frac{E_y^{\mathrm{miss}}}{E_x^{\mathrm{miss}}}\right).$$

The vector $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ thus represents the total undetected transverse momentum in the event via its magnitude $E_{\mathrm{T}}^{\mathrm{miss}}$, and its direction in the transverse plane is defined in terms of ϕ^{miss}. In addition to genuine contributions from invisible particles, false signatures of missing transverse momentum can arise from detector effects such as finite resolution, limited acceptance, miscalibration of physics objects, detector noise, and inactive detector regions. It is therefore essential to account for these effects in order to obtain a well-reconstructed $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$.

6.6 Overlap removal

Even after the reconstruction of individual physics objects using the procedures described above, it is possible for the same energy deposit or track in the detector to be used in the reconstruction of multiple objects in different detector subsystems. Such ambiguities arise primarily due to the finite detector resolution and overlapping signatures of particles, such as when a lepton is reconstructed in close proximity to a jet. To mitigate this problem and ensure a well-defined set of physics objects, an overlap removal procedure is applied.

For the analyses presented in this thesis, overlap removal follows the standard ATLAS recommendation, summarised in Table 6.3. The separation between two reconstructed objects is quantified by $\Delta R = \sqrt{\Delta y^2 + \Delta\phi^2}$. The overlap removal procedure is applied sequentially. First, electrons that share an ID track with a muon are removed, since such electrons typically originate from Bremsstrahlung photons that are emitted. Next, among the reconstructed small-R jets, those found within $\Delta R(\text{electron}, \text{jet}) < 0.2$ are discarded and assumed to have originated from the electron. Then, if an electron is found close to a remaining jet within $\Delta R(\text{electron}, \text{jet}) < 0.4$, the electron is removed, as it is likely to have originated from jet fragmentation rather than from an isolated electron. Jets associated with fewer than three tracks are then removed if they are within $\Delta R < 0.2$ from a muon candidate. Finally, muons within the range $0.2 < \Delta R < 0.4$ of the remaining jets are removed to suppress non-prompt muons.

Table 6.3: Table illustrating the criteria used for the various steps of the overlap removal procedure.

Reject	**Against**	**Criteria**
Electron	Muon	Shared ID track
Jet	Electron	$\Delta R < 0.2$
Electron	Jet	$0.2 < \Delta R < 0.4$
Muon	Jet	$\Delta R < 0.4$

CHAPTER 7

Probing entanglement close to the $m_{t\bar{t}}$ threshold

"What we know is a drop, what we don't know is an ocean."

— Isaac Newton

Building on the theoretical foundations presented in Chapter 3, in this chapter the experimental measurement of quantum entanglement in $t\bar{t}$ pairs produced near the $m_{t\bar{t}}$ threshold is described, and the corresponding results are presented. The measurement is performed in the lepton+jets final state and builds upon previous work as well as insights gained from the corresponding measurement in the dilepton channel [17]. In contrast to the dilepton analysis, contributions from $t\bar{t}$ quasi-bound-state effects, discussed in Section 3.3, are also taken into account near the $t\bar{t}$ production threshold. Therefore, the signal definition includes both pQCD $t\bar{t}$ production and contributions from $t\bar{t}_{\text{GFRW}}$, making this the first ATLAS entanglement measurement to incorporate such effects. Results obtained using the simplified η_t model are also presented for comparison.

As discussed in Section 3.2, the measurement relies on the accurate identification of spin analysers from the top- and antitop-quark decay products and on reconstructing the angle between them in their respective top-quark rest frames. In this measurement, the ***charged lepton*** serves as the spin analyser of the semi-leptonically decaying top quark, while the ***down-type jet*** from the hadronic W-boson decay is used as the hadronic spin analyser. Identifying the down-type jet is particularly challenging due to similar kinematic properties of the two jets originating from the hadronic W-boson decay. Its identification therefore relies on indirect methods, including flavour-tagging information and additional kinematic properties that can be used to assign jets to the corresponding top-quark decay partons during the reconstruction of the $t\bar{t}$ system.

7.1 Analysis strategy

The goal of the measurement is to extract the entanglement marker D from the mean of the $\cos\varphi$ distribution (Equation 3.10). Once determined, it is verified whether under the assumption of QM within the SM, the $t\bar{t}$ system is entangled near the $m_{t\bar{t}}$ threshold. This is achieved by measuring the degree of entanglement in data and comparing the result to MC predictions, which provide the best available description of the entanglement expected in the SM. Performing this measurement requires

several analysis steps, and the workflow proceeds as follows: Following the reconstruction of physics objects described in Chapter 6, baseline objects satisfying specific requirements such as isolation and identification are selected (Section 7.2). To maximise the purity of $t\bar{t}$ events in the lepton+jets channel, a dedicated event selection is then applied (Section 7.3). The $t\bar{t}$ system is then fully reconstructed using an attention-based machine learning algorithm, SPANet (Section 7.4). In addition, an alternative reconstruction based on a χ^2 method is employed at particle level (Section 7.5).

The reconstructed event kinematics are further subjected to selections designed to define a signal region sensitive to entanglement effects while enhancing the purity of events in which the down-type jet is correctly identified (Section 7.6). Finally, the reconstructed angular distributions are used to extract the entanglement-sensitive observable, D, which is corrected to its true value using a calibration curve method described in Section 7.7. The construction of this curve also incorporates the various sources of uncertainty affecting the measurement, as described in Section 7.9. Background contributions arising from non-prompt and fake leptons are estimated using a data-driven method described in Section 7.8. The kinematic distributions entering the measurement are discussed in Section 7.10. The measurement results and their interpretation are presented in Section 7.11, where they are also compared to a scenario in which $t\bar{t}$ quasi-bound-state effects are not included.

7.2 Baseline object selection

7.2.1 Leptons

Selected electrons are required to satisfy the *Tight LH* identification WP (Section 7.2.1). To ensure isolation, the *FCTight* isolation WP is applied. To select electrons originating from the primary vertex, the transverse impact parameter $|d_0|$ is required to be smaller than 5 times the uncertainty in its measurement (σ_{d_0}), while the longitudinal impact parameter must fulfill $|z_0 \sin\theta| < 0.5$ mm. Furthermore, electrons are required to have a transverse momentum $p_T > 10$ GeV and the energy cluster needs to fall in the pseudorapidity range $|\eta_{\text{cluster}}| < 2.47$, excluding the transition region between the barrel and endcap calorimeters ($1.37 < |\eta_{\text{cluster}}| < 1.52$). This transition region contains large amounts of uninstrumented material, which degrades the electron energy measurement and therefore worsens the energy resolution. A looser electron selection is used for the definition of control regions employed in the data-driven estimation of non-prompt or misidentified ("fake") lepton backgrounds (Section 7.8).

Muons follow a selection similar to that of electrons. For identification, the *Medium* quality criteria are applied, while isolation is based on the *FCTightTrackOnly* WP. Muon candidates are required to have a transverse momentum $p_{\text{T}} > 10$ GeV and a pseudorapidity $|\eta| < 2.5$. In addition, requirements on the transverse impact parameter significance and the longitudinal impact parameter are imposed to ensure that the muons originate from the primary vertex. Similar to electrons, a looser muon selection is also used for the definition of "fake" control regions. The requirements for the electron and muon selection are summarised in Table 7.1.

7.2.2 Small-radius jets

All small-R jets are reconstructed using the anti-k_T jet clustering algorithm with a radius parameter of $R = 0.4$, using PFlow objects as inputs (Section 6.4.3). Jet candidates are required to satisfy $p_{\text{T}} > 25$ GeV and $|\eta| < 2.5$. To further mitigate pile-up effects, a JVT requirement (Section 6.4.3) is

Table 7.1: Selection criteria for electrons and muons. The looser criteria applied in the definition of the "fake" control regions are shown in italics and square brackets.

Feature	**Electron**	**Muon**
Pseudorapidity range	$\lvert\eta_{\text{cluster}}\rvert < 1.37$ or $1.52 < \lvert\eta_{\text{cluster}}\rvert < 2.47$	$\lvert\eta\rvert < 2.5$
Transverse momentum	$p_T > 10$ GeV	
Track to vertex association	$\lvert d_0/\sigma_{d_0}\rvert < 5$ $\lvert z_0 \sin\theta\rvert < 0.5$ mm	$\lvert d_0/\sigma_{d_0}\rvert < 3$ $\lvert z_0 \sin\theta\rvert < 0.5$ mm
Identification	Tight LH [*Medium LH*]	Medium [*Loose*]
Isolation	FCTight [*no isolation*]	FCTightTrackOnly [*no isolation*]

applied to jets with $p_{\text{T}} < 60$ GeV and $|\eta| < 2.4$. The jet selection requirements are summarised in Table 7.2.

Table 7.2: Selection criteria for small-R jets.

Feature	**Criteria**
Pseudorapidity range	$\lvert\eta\rvert < 2.5$
Transverse momentum	$p_T > 25$ GeV
Pile-up suppression	JVT> 0.50 for $p_T < 60$ GeV and $\lvert\eta\rvert < 2.4$

7.3 Event selection

The measurement focuses on $t\bar{t}$ events in final states containing a single electron or muon, missing transverse momentum, and jets. Initially, general data-quality requirements, referred to as preselections, are applied. Furthermore, events are required to satisfy a dedicated set of single-lepton triggers, ensuring that their detector signatures are consistent with the targeted final state. Finally, object-related selection criteria are applied to the reconstructed physics objects in order to retain events consistent with $t\bar{t}$ decays in the lepton+jets channel.

7.3.1 Preselection

Only events satistfying a set of general quality requirements are retained for further analysis. Requirements are imposed to reject events affected by temporary subdetector malfunctions, such as noise bursts in the LAr calorimeter, temporary failures or recovery procedures in the SCT tracking detector caused by radiation-induced electronics upsets, corrupted data from the Tile calorimeter, or incomplete detector information. In addition, events are required to belong to a luminosity block contained in the Good Runs List (Section 4.2.6), during which all detector subsystems are fully functional and there are stable beams.

Only events containing a primary vertex with at least two associated tracks are retained. The primary vertex is defined as the one with the largest $\sum p_T^2$, where the sum runs over all tracks with

transverse momentum $p_T > 0.5$ GeV associated with the vertex. This requirement helps suppress events arising from external sources, such as cosmic rays or background radiation in the cavern.

7.3.2 Trigger selection

To preselect events potentially associated with the signal of interest, events are required to satisfy a dedicated set of HLT triggers (Section 4.2.5) requiring the presence of at least one lepton in the final state. The same trigger selection is applied to the MC samples, together with correction factors that account for differences in trigger efficiencies between data and simulation.

For this analysis, the unprescaled single-electron [225] and single-muon [226] triggers, with the lowest available transverse momentum (p_T) thresholds, are used. The threshold depends on the data-taking year. To maintain acceptable trigger rates while allowing for lower p_T thresholds, these triggers include an additional lepton-isolation requirement, which suppresses backgrounds from non-prompt leptons. However, the efficiency of isolated triggers can decrease in events with large hadronic activity, where genuine leptons from $t\bar{t}$ events may fail the isolation requirement because the decay products become more collimated and the lepton can be located close to jets. To recover such events, the isolated triggers are combined with non-isolated single-lepton triggers but with higher p_T thresholds, which compensate for the larger rates of non-prompt, non-isolated leptons passing the trigger selection. This combination ensures a stable trigger efficiency across the full $t\bar{t}$ invariant mass spectrum. Table 7.3 summarises the full set of triggers and their corresponding requirements. Events are required to satisfy at least one of these triggers, implemented as a logical OR of the selected single-electron and single-muon triggers.

Table 7.3: Summary of the single-lepton triggers used in this analysis and their requirements.

Year	p_T threshold (GeV)	Isolation
Electron triggers		
2015	24	Yes
2015	60	No
2015	120	No
2016-2018	26	Yes
2016-2018	60	No
2016-2018	140	No
Muon triggers		
2015	20	Yes
2015	50	No
2016-2018	26	Yes
2016-2018	50	No

7.3.3 Common object-related selection

Since the analysis targets $t\bar{t}$ decays in the lepton+jets channel, events are required to contain exactly one electron or muon that passes the object selection criteria described in Section 7.2. The selected lepton must have transverse momentum, $p_T > 28$ GeV. This threshold is chosen to lie above all of the trigger

thresholds for isolated leptons, ensuring full trigger efficiency. In addition, the lepton is required to be geometrically matched to the corresponding trigger object using an angular separation criterion of $\Delta R < 0.15$, ensuring that the reconstructed lepton is consistent with the object that satisfies the trigger requirement, thereby providing a reliable correspondence between trigger-level and reconstructed objects. Events containing any additional lepton with $p_T > 10$ GeV are vetoed in order to suppress backgrounds with multiple prompt leptons, such as dileptonic $t\bar{t}$ and diboson production.

Events must also include at least four jets satisfying the requirements detailed in Section 7.2. Among these, exactly two jets must be b-tagged using the DL1r algorithm at 60% efficiency WP (Section 6.4.6). Events are also accepted if they contain one additional jet, distinct from these two, that is b-tagged at the 85% efficiency WP provided it fails the 60% WP requirement. This requirement allows jets compatible with charm or light-flavour hadrons to enter the selection, as they are most likely to satisfy the 85% efficiency WP while failing the tighter 60% selection.

To further suppress backgrounds from processes involving non-prompt or fake leptons, and to account for the presence of neutrinos in the event, a requirement of missing transverse momentum $E_{\mathrm{T}}^{\mathrm{miss}} > 20$ GeV is applied. A summary of all the event selection criteria is provided in Table 7.4.

These event selections define the signal region used for training the SPANet algorithm for event reconstruction at detector level, as described in Section 7.4. Starting from this signal region definition, additional selection criteria are applied after $t\bar{t}$ reconstruction to further tighten the selection for the final analysis signal region (Section 7.6).

Table 7.4: Summary of the event selection criteria.

Selection	**Criteria**
Data quality criteria	Data events in GRLs (2015-2018); jet cleaning, primary vertex selection, removal of LAr noise bursts, tile and SCT errors.
Single lepton trigger	Separate single-electron or single-muon triggers
Exactly one lepton	== 1 e or μ with $p_T > 28$ GeV No additional lepton with $p_T > 10$ GeV.
$E_{\mathrm{T}}^{\mathrm{miss}}$ cut	$E_{\mathrm{T}}^{\mathrm{miss}} > 20$ GeV
At least four jets	≥ 4 jets, $p_T > 25$ GeV
b-tagging	Exactly 2 b-tagged jets at 60% WP and at most one b-tagged jet at 85% WP

7.4 Reconstruction of the $t\bar{t}$ system

Once the required physics objects are selected, the next step is the full reconstruction of the $t\bar{t}$ system. Since this measurement is particularly sensitive to the correct identification of the down-type quark jet, the correct assignment of reconstructed objects to their corresponding spin analysers is a key challenge. This requirement goes beyond the needs of many $t\bar{t}$ searches and measurements, such as top-quark mass measurements or analyses using the $t\bar{t}$ invariant mass spectrum, where it is sufficient to identify the correct set of jets originating from the hadronic top-quark decay, without requiring a precise assignment to the underlying partons. To address this assignment challenge, a machine learning algorithm known as SPANet [227–229] is used. It is based on a transformer architecture and employs an attention mechanism designed to tackle the combinatorial problem of identifying and assigning

jets originating from top-quark decays. In particular, it matches detector-level reconstructed jets to the underlying partons that initiated the hadronic activity, even in events with high jet multiplicities. Unlike traditional approaches, SPANet evaluates all possible permutations simultaneously and assigns a score to each configuration, enabling the identification of the most probable assignment. This strategy avoids the combinatorial explosion encountered in methods such as KLFitter [230] or fully connected neural networks [231].

At particle level, the $t\bar{t}$ system is reconstructed using a χ^2 minimisation algorithm (Section 7.5). While it would have been desirable to use SPANet at particle level to ensure consistency in the reconstruction procedure across both detector and particle levels, this was not possible. SPANet was trained exclusively on detector-level events, and retraining it for particle-level events was not feasible within the timeline of the analysis. Therefore, the χ^2 algorithm is used as the most suitable alternative for particle-level reconstruction.

7.4.1 SPANet

SPANet, short for ***Symmetry Preserving Attention Network***, is a deep learning model designed to address the combinatorial challenge of identifying jets originating from top-quark decays. The network is built on a transformer architecture [227], as illustrated in Figure 7.1. This architecture offers the ability to handle variable-length inputs and produces outputs that are equivariant under permutations of the input ordering. Both of these properties are well-suited to $t\bar{t}$ analyses, where the number of reconstructed jets can vary between events and there is no inherent ordering of jets.

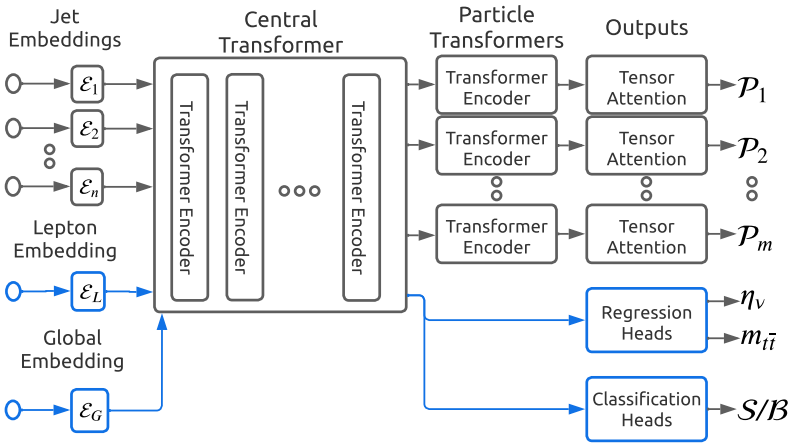


Figure 7.1: Diagram of the SPANet architecture. Arranged from left to right, the inputs are denoted by ϵ_i, the assignment outputs by $\mathcal{P}_1 \ldots \mathcal{P}_m$, the regression outputs by η_ν and $m_{t\bar{t}}$, and the classification output by $\mathcal{S}/\mathcal{B}$ [229]. In this measurement, $m = 2$, corresponding to the two top-quark assignment probability distributions.

In SPANet, the kinematic input features of jets $(\epsilon_1, \ldots, \epsilon_n)$ and leptons (ϵ_L), along with some global features (ϵ_G) such as information on missing transverse momentum, are embedded into a high-dimensional latent space. These embeddings are then processed by a central *transformer encoder*, which provides contextual information by modelling correlations between all input objects. Following this, separately trained *particle transformers*, one for each top quark, are applied to predict

the assignment of jets to individual top-quark decay products. The original SPANet formulation, developed for observables such as $m_{t\bar{t}}$ and top-quark mass reconstruction, is based on a key symmetry assumption: the decay products originating from the W-boson decay of the hadronic top quark are treated as permutation invariant. In those applications, it is not necessary to distinguish between the jets originating from the up- and down-type quarks, and this symmetry is incorporated into the network to improve performance by reducing the combinatorial complexity. In the present measurement, however, this symmetry assumption is explicitly removed, since the correct assignment of the down-type quark jet is essential for the entanglement measurement. This is achieved by modifying the attention mechanism, in particular by removing the permutation-symmetry assumption in Equation 2 of Ref. [227], allowing the network to learn ordered assignments for the W-boson decay products.

To obtain a consistent assignment across all particles in an event, a symmetric *tensor attention* layer is employed in each particle transformer. This layer extends the standard attention mechanism by accounting for residual permutation symmetries while incorporating the aforementioned modification that breaks the W-boson decay symmetry within each transformer, thereby enabling ordered treatment of the decay products where required. The output is a joint probability distribution describing the likelihood of a given jet combination corresponding to a specific top quark, denoted by $\mathcal{P}_1$ and $\mathcal{P}_2$ (see Figure 7.1). The most likely jet assignment is then taken from the peak of each distribution. In cases where the assignments overlap, the assignment with the higher probability is selected first, and the second distribution is re-evaluated to determine the best non-conflicting assignment.

In addition to jet assignment, SPANet can also be extended to perform auxiliary tasks such as signal vs background classification of events and kinematic regression of kinematic variables (shown in blue on the right of Figure 7.1), for example estimating the longitudinal component of the neutrino momentum, by utilising the optimised latent space learned in the reconstruction task.

As a supervised learning model, SPANet requires dedicated training on simulated $t\bar{t}$ events with known parton-level information. This training step is essential for the network to learn how to distinguish between correct and incorrect jet-parton assignments based on the reconstructed event kinematics.

Training and performance

The SPANet algorithm takes as input the four-momenta of the reconstructed objects, along with the missing transverse momentum and its direction. For jets, the PCBT score (Section 6.4.6) is included, while for leptons two binary flags are used to distinguish between electrons and muons. To avoid ambiguities arising from the periodicity of the azimuthal angle ϕ at $\pm\pi$, ϕ is encoded using its sine and cosine components, $\sin\phi$ and $\cos\phi$. All input features with energy or momentum units are provided to the network in logarithmic form to reduce variance and improve numerical stability during training. Target assignments are defined by matching partons from the MC record to reconstructed jets using a distance criterion of $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} < 0.3$. A semi-leptonically decaying top quark is considered correctly reconstructed if the b-quark from its decay is matched to a jet. For the hadronically decaying top quark, a successful parton matching requires that all three decay partons, the b-quark, the up-type quark, and the down-type quark, are each matched to distinct jets. In cases where multiple partons are matched to the same jet, only the assignment with the smallest ΔR is retained. This matching procedure defines the training targets for the network.

The SPANet configuration used in this analysis is trained on the nominal POWHEG+PYTHIA $t\bar{t}$ MC sample. This training is performed after applying the basic event selection described in Section 7.3, requiring events with at least four jets, of which exactly two are b-tagged using the 60% WP requirement. The dataset is split into two equally sized subsets, and two networks are trained in a *cross-training* scheme: each network is trained on one subset and evaluated on the other. Approximately 5% of each training subset is reserved for validation.

SPANet contains a number of *hyperparameters* that define both the network architecture and the training procedure. Hyperparameter optimisation refers to the process of selecting the best hyperparameters for a machine learning model for a given task. Several strategies exist for this, ranging from simple methods such as grid or random search to more advanced approaches. In this analysis, a dedicated hyperparameter optimisation is performed using the built-in SPANet tuning framework. Parameters such as the hidden dimension, the number of layers in each network block, and the learning rate, are optimised. The optimisation is carried out using a parallel Gaussian process search implemented in the SHERPA hyperparameter optimisation library [232].

Compared to the original SPANet training, the training procedure used in this analysis is modified to account for the correct assignment of the up- and down-type jets originating from the hadronic W-boson decay. This additional distinction increases the complexity of the jet-assignment task and results in slightly longer training times. After training, events are reconstructed using the trained networks. The reconstruction performance is evaluated by comparing the assigned jets to their true partonic origins, based on the truth-matching procedure described above. The reconstruction efficiency, defined as the fraction of correctly matched or well-reconstructed events among all parton-matched events, is found to be 54.1% in the inclusive $m_{t\bar{t}}$ region. Further optimisation using selections on SPANet performance variables (Section 7.6) leads to an increased purity of well-reconstructed events.

Performance variables

In addition to the jet-parton assignment, SPANet produces several auxilliary outputs that help quantify the quality of the event reconstruction. For each reconstructed top quark, SPANet provides two key probabilities. The ***assignment probabilty*** reflects the likelihood that the selected jets correspond to the truth-matched partons, assuming the target particle is reconstructable. The ***detection probability***, on the other hand, represents the likelihood that the top quark is reconstructable at all, that is, each of its decay products can be matched to a reconstructed jet. These probabilities are derived by adding two additional output heads to the central transformer block of the SPANet architecture. The assignment probability is estimated by adding a learnable event vector to the transformer inputs. This vector captures global event information and is processed by the central transformer. Its output is then passed through a feed-forward neural network to predict the assignment probability. The detection probability is computed in a similar way. However, instead of using the event vector from the central transformer, it uses a summary particle vector obtained from each particle transformer encoder [229].

By applying optimised selection criteria on these assignment and detection probabilities, one can enhance the fraction of correctly reconstructed events included in this analysis. This, in turn, helps reduce background processes not involving top quarks as well as reject $t\bar{t}$ events where one or more of the top quark decay products are not detected. The specific selections applied on these variables are described in detail in Section 7.6.

7.4.2 Kinematic reconstruction of the $t\bar{t}$ system

Once the b-, up-type, and down-type jets are assigned by SPANet, the $t\bar{t}$ system can be reconstructed. However, for a complete reconstruction of the $t\bar{t}$ system, determining the contribution from the neutrino coming from the semi-leptonically decaying top quark is essential. Since the neutrino is one of those physics objects that is impossible to directly detect, its kinematic reconstruction is not straightforward. The only available information about the neutrino is encoded in the form of the missing transverse momentum, characterised by its magnitude, E_T^{miss}, and its direction in the transverse plane, ϕ^{miss} (Section 6.5). While the transverse components of the neutrino momentum, p_x^ν and p_y^ν, can be inferred from E_T^{miss} and ϕ^{miss}, the longitudinal momentum component, p_z^ν, cannot be measured directly and therefore remains unknown. This component can either be determined by SPANet via regression (Section 7.4) or determined analytically by imposing a mass constraint on the W boson. In this analysis, both approaches have been found to yield similar results, and therefore the traditional analytical reconstruction approach is used. In this method, the transverse momentum of the neutrino $\vec{p}_{\text{T}}^{\,\nu} = (p_x^\nu,\ p_y^\nu)$ is taken to be the missing transverse momentum, $\vec{E}_{\text{T}}^{\text{miss}}$. The longitudinal component, p_z^ν, is then calculated using a kinematic constraint. Assuming that the W boson from the semi-leptonically decaying top quark is on-shell, the invariant mass of the lepton-neutrino system satisfies:

$$(p^\nu + p^l)^2 = m_W^2$$

where p^ν and p^l denote the four-momenta of the neutrino and the charged lepton, respectively. The W-boson mass is taken as $m_W = 80.4$ GeV [31]. This condition leads to a quadratic equation in p_z^ν:

$$a(p_z^\nu)^2 + bp_z^\nu + c = 0, \tag{7.1}$$

where,

$$\begin{aligned} a &= (E^l)^2 - (p_z^l)^2, \\ b &= -2kp_z^l, \\ c &= (E^l)^2(p_T^\nu)^2 - k^2, \\ k &= \frac{(m_W^2 - m_l^2)}{2} + (p_x^l p_x^\nu + p_y^l p_y^\nu). \end{aligned}$$

Here, E^l and m_l denote the energy and mass of the charged lepton. Equation 7.1 is solved for p_z^ν. If exactly one real solution exists, it is used directly. If no real solution exists (which occurs in approximately one-third of events), the real part, $p_z^\nu = -b/(2a)$, is taken. If two real solutions are found, the solution yielding a reconstructed top-quark mass closest to 172.5 GeV is selected.

The four-momenta of the hadronically and semi-leptonically decaying top quarks are then reconstructed from their decay products. The four-momentum of the hadronically decaying top quark, t_{had}, is formed from the four-momenta of the jets assigned to the corresponding bottom quark and the two light quarks from the hadronic W-boson decay. Similarly, the four-momentum of the semi-leptonically decaying top quark, t_{lep}, is reconstructed from the four-momenta of the jet assigned to the corresponding

bottom quark, the charged lepton, and the neutrino. This can be expressed as:

$$\left(p_{\ell^-} + p_{\bar{\nu}}\right)^2 = m^2_{W_{\text{lep}}}, \left(p_u + p_{\bar{d}}\right)^2 = m^2_{W_{\text{had}}}.$$
$$\left(p_{\ell^-} + p_{\bar{\nu}} + p_{\bar{b}}\right)^2 = m^2_{t_{\text{lep}}}, \left(p_u + p_{\bar{d}} + p_{\bar{b}}\right)^2 = m^2_{t_{\text{had}}}.$$

The invariant mass of the $t\bar{t}$ system, $m_{t\bar{t}}$, which is crucial for the definition of the threshold region in the entanglement measurement, is then calculated using the reconstructed four-momenta of the hadronically and semi-leptonically decaying top quarks.

Transforming to the parent top-quark rest frames

The angular observable $\cos\varphi$, used for the measurement of entanglement, is constructed from the directions of the charged lepton from the semi-leptonically decay top-quark and the down-type jet coming from the W boson of the hadronically decaying top quark. This angle is measured by taking the directions of these two spin analysers in their respective top-quark rest frames. By default, the four-momenta of the charged lepton and the down-type jet are defined in the laboratory frame. From Section 3.2, it is clear that spin correlations are defined well in the $t\bar{t}$ centre-of-mass frame where the spin quantisation axes are defined consistently. Therefore, to transform the directions of the spin analysers from the laboratory to the respective top-quark rest frames, a consistent sequence of Lorentz transformations must be applied.

First, the four-momentum of the reconstructed $t\bar{t}$ system is obtained in the laboratory frame. A Lorentz boost is then performed to transform all relevant four-momenta into the $t\bar{t}$ centre-of-mass frame. In this frame, the top and anti-top quarks are back-to-back, providing a well-defined reference for spin correlation measurements. Subsequently, separate Lorentz boosts are applied to each top-quark decay system. The decay products of the semi-leptonically decaying top quark are boosted into the rest frame of the parent top quark, and similarly, the decay products of the hadronically decaying top quark are boosted into its corresponding rest frame. These boosts are constructed from the reconstructed four-momenta of the respective top quarks. Once in their respective parent top-quark rest frames, the three-momentum directions of the charged lepton and the down-type jet are extracted and normalised. The observable $\cos\varphi$ is then defined as the scalar product of these unit vectors. This two-step procedure ensures that both decay systems are defined with respect to a common reference frame. When changing between reference frames, the combination of Lorentz boosts can induce a rotation of the spatial coordinate axes, known as a Wigner rotation [233]. In the case of a direct boost from the laboratory frame to the individual top-quark rest frames, this leads to an incorrect Wigner rotation of the coordinate axes. As a result, the directions of the spin analysers would be defined in misaligned coordinate systems and the resulting $\cos\varphi$ would no longer correspond to a well-defined physical observable. At the particle and the parton level, $\cos\varphi$ is defined in an analogous manner.

7.5 Particle-level selection and reconstruction of the $t\bar{t}$ system

Particle-level objects are reconstructed from stable particles in the simulation, where stable particles are defined as those with a mean lifetime greater than 30 ps (Section 5.3). Only electrons, muons, and neutrinos originating from W- or Z-boson decays, including via intermediate τ-lepton decays, are considered. Leptons originating from the decay of hadrons or a τ-lepton are identified using the MC

truth-record, which stores the simulated generator-level particle history, and are excluded from the selection. Final-state photons, γ, with a cone of $\Delta R(l,\gamma) < 0.1$ around charged leptons are combined with the lepton four-momentum in a procedure known as *dressing*. Electrons and muons are required to satisfy $p_T > 10$ GeV and $|\eta| < 2.5$. Neutrinos are not dressed and are not subject to additional phase-space requirements.

Jets are defined using the same anti-k_T algorithm as used at the reconstruction level with a radius paramater of $R = 0.4$. For the jet clustering, all stable particles are considered except for the electrons, muons, neutrinos, and photons used in the lepton definitions. The resultant particle-level jets are required to satisfy $p_T > 25$ GeV and $|\eta| < 2.5$. The identification of jet flavour at particle level is commonly performed using a *ghost-matching* procedure [234]. In this method, flavour-carrying particles, typically hadrons, are included in the jet clustering as "ghost" particles by rescaling their four-momenta to negligible values. This ensures that they do not affect the clustering outcome or the jet kinematics, while still allowing them to be clustered into jets. The flavour of a jet can then be defined through the presence of such ghost particles within the jet. In this analysis, b-jets are identified using this procedure. A jet is labelled as b-tagged if it contains one ghost-associated B-hadron with $p_T > 5$ GeV.

As at the detector level, the analysis relies on the correct assignment of the four jets originating from the $t\bar{t}$ decay, particularly the down-type jet. To achieve this optimal assignment, a ***χ^2 minimisation technique*** is employed, following approaches adopted in various other $t\bar{t}$ analyses [25, 235]. The χ^2 function to be minimised is defined as follows:

$$\chi^2 = \left[\frac{m_{jj} - m_W}{\sigma_W}\right]^2 + \left[\frac{(m_{jjb} - m_{jj}) - m_{(t_h - W)}}{\sigma_{(t_h - W)}}\right]^2 + \left[\frac{m_{jl\nu} - m_{t_l}}{\sigma_{t_l}}\right]^2 + \left[\frac{(p_{T,jjb} - p_{T,jl\nu}) - (p_{T,t_h} - p_{T,t_l})}{\sigma_{\Delta p_T}}\right]^2 . \tag{7.2}$$

The first term in the function constrains the invariant mass of a pair of jets, m_{jj}, to be close to the mass of the W boson. The second term requires the invariant mass of a three-jet system, m_{jjb}, to be close to the mass of the hadronically-decaying top quark, m_{t_h}. However, since m_{jj} and m_{jjb} are strongly correlated, the mass of the hadronically-decaying W boson is subtracted from m_{t_h}, resulting in a decorrelated term, $m_{(t_h - W)}$. The third term is analogous to the second but applies to the semi-leptonically decaying top quark, m_{t_l}. The last term imposes the constraint that the momenta of the two reconstructed top quarks are similar, reflecting momentum balance in the $t\bar{t}$ system. The values of the parameters in the χ^2 function are derived from pQCD $t\bar{t}$ events at particle level, where the lepton, reconstructed neutrino, and the three jets are matched to the correct permutation using the truth information from the generator. All parameters are extracted from Gaussian fits on kinematic distributions, provided in Appendix A.

All possible jet permutations, including all jets that pass the selection requirements described in Section 7.3, are considered, and the combination yielding the lowest χ^2 value is selected. The four momentum of the neutrino is reconstructed analytically using the W-boson mass constraint. In cases where the neutrino longitudinal momentum has two real solutions, the solution that yields the lowest χ^2 is chosen. To reduce the number of permutations, only jets identified as b-jets via ghost matching

are allowed to be assigned to either of the b-quarks from the top quark decays, thereby reducing the number of possible permutations.

To evaluate the performance of this χ^2 minimisation technique, the fraction of particle-level events with correctly assigned jets is determined by matching the selected jets to their corresponding parton-level objects within a cone of $\Delta R < 0.3$. This study is performed using pQCD $t\bar{t}$ events, where approximately 55% of events in the low $m_{t\bar{t}}$ region are found to be have correct jet assignments. This fraction can further be improved by applying an additional requirement on $\log_{10}(\chi^2)$, which can be optimised to select only well-reconstructed events at the cost of reduced statistics. Table 7.5 summarises the impact of this requirement on the assignment efficiency of events in the low $m_{t\bar{t}}$ region. As expected, the efficiency increases with tighter selections on $\log_{10}(\chi^2)$, demonstrating the effectiveness of the χ^2 minimisation method in correctly assigning jets to the respective top quarks. Despite the $\log_{10}(\chi^2)$ selection being useful, no such selection is applied in this analysis. This choice ensures maximum event retention at particle level and helps avoid migrations between $\cos\varphi$ distribution bins at particle- and detector-levels during the calibration curve construction (Section 7.7).

Table 7.5: Summary of the fraction of correctly assigned jets for different selection requirements on the value of $\log_{10}(\chi^2)$, following the application of the χ^2 minimisation technique to determine the optimal combination of jets at the particle level.

Selection	χ^2 jet assignment efficiency
no $\log_{10}(\chi^2)$ cut	55.2 %
$\log_{10}(\chi^2) < 1.3$	76.2 %
$\log_{10}(\chi^2) < 1.2$	78.4 %
$\log_{10}(\chi^2) < 1.1$	80.4 %
$\log_{10}(\chi^2) < 0.9$	83.7 %

7.6 Signal and validation region definitions

Since entanglement effects are expected to be most prominent near the $t\bar{t}$ production threshold, additional selection criteria are applied to define the final ***signal region*** of the analysis, thereby enhancing the purity of entangled $t\bar{t}$ states. To validate the measurement strategy in regions of phase space different from the signal region, two additional regions, referred to as ***validation regions***, are also defined.

7.6.1 Reconstruction level

As mentioned previously, the sensitivity to entanglement depends strongly on the selected $m_{t\bar{t}}$ phase space, the accurate identification of the spin analysers originating from the top- and antitop-quark decays, as well as on the reconstruction of the angle between them in their respective parent top-quark rest frames. While the construction of this angle has already been described in Section 7.4.2, and the assignment of reconstructed jets to the corresponding decay products is performed using SPANet (Section 7.4.1), this section focuses on the determination of spin analysers entering the angle definition,

optimised to enhance the sensitivity to entanglement, together with additional kinematic selections serving the same purpose.

The charged lepton from the semi-leptonically decaying top quark can be identified with high purity due to its clean experimental signature. The identification of the down-type jet from the hadronic W-boson decay, however, is significantly more challenging, since no dedicated tagging algorithm currently exists for distinguishing jets originating from down-type quarks in a manner analogous to the b-tagging algorithms available for b-jets. Therefore, the identification of the down-type jet must rely on indirect flavour-sensitive information.

As discussed in Section 6.4.6, the DL1r b-tagging discriminant exhibits a region enriched in c-jets. In the context of hadronic W-boson decays, this region enhances sensitivity to $W \to cs$ and $W \to cd$ processes, which account for approximately half of all hadronic W-boson decays and thus provide the dominant source of c-jets. Selecting c-jet-enriched events therefore increases the probability of correctly identifying the corresponding down-type jet. This provides a practical way of exploiting flavour information to improve the assignment of the hadronic spin analyser in these events. To exploit this feature, the event selection described in Section 7.3.3 allows events containing an additional jet passing the 85% efficiency WP of the DL1r b-tagging algorithm while failing the 60% requirement. Such events are expected to have an enhanced fraction of c-jets.

The events are subsequently processed by SPANet (Section 7.4.1), which provides an ordered assignment of reconstructed jets to the corresponding top- and antitop-quark decay products. Following this assignment, additional flavour-sensitive information based on the PCBT scores is exploited in order to enhance the fraction of correctly assigned down-type jets in charm-enriched event topologies. In particular, the two jets assigned by SPANet to the hadronic W-boson decay are required to have PCBT scores below 5 (Section 6.4.6), ensuring that neither jet is likely to originate from b-quarks. In addition, these two jets must have different PCBT scores so that the jets are distinguishable. The jet with the lower score is then chosen as the down-type jet. Events failing this criterion are discarded. This selection exploits the fact that in $W \to cs$ and $W \to cd$ decays the c-quark typically produces a jet with a higher b-tagging score than the accompanying light-flavour jet, due to its longer lifetime and more pronounced secondary-vertex structure. This leads to a separation in the b-tagging score distributions of the two jets, enabling a more reliable identification of the down-type jet.

Events originating from $W \to ud$ and $W \to us$ decays, where both jets are light-flavoured, may also enter this selection, although their contribution is expected to be suppressed by the event selection requirements. In these decays the ordering of jets based on the b-tagging score becomes effectively random and therefore does not provide additional discrimination power. Overall, this procedure used to determine the down-type jet spin analyser in a charm-enriched region can be viewed as an *indirect c-tagging* approach based on b-tagging information.

To further improve the fraction of correctly reconstructed events and suppress those that are poorly reconstructed or non-reconstructable, several selection requirements on the SPANet performance variables are tested. Given the strong dependence of the analysis on correct down-type jet identification, a requirement on the hadronic top detection probability, `spanet_hadtop_detection` > 0.7, is found to be optimal based on the expected sensitivity to entanglement. Table 7.6 summarises the fractions of correctly matched, incorrectly matched, and non-reconstructable down-type jets before and after this selection requirement. Here, correctly matched events correspond to those in which the reconstructed down-type jet is associated with the true parton-level down-type quark. Incorrectly matched events include cases of misassignment, while non-reconstructable events are those in which no reconstructed jet can be matched to the parton-level down-type quark. As illustrated in Figure 7.2,

the fraction of correctly matched down-type jets increases significantly with increasing values of `spanet_hadtop_detection`, demonstrating the effectiveness of this requirement.

To isolate events near the $t\bar{t}$ production threshold, a requirement on the invariant mass of the $t\bar{t}$ system to satisfy $m_{t\bar{t}} < 390\,\text{GeV}$, is made. To further enrich the gg fusion contribution, which dominates the production of entangled $t\bar{t}$ pairs in this threshold region, an additional requirement is imposed on the boost of the $t\bar{t}$ system, $\beta_{t\bar{t}} = p_{t\bar{t}}/E_{t\bar{t}} < 0.55$, where $p_{t\bar{t}}$ denotes the magnitude of the three-momentum of the $t\bar{t}$ system and $E_{t\bar{t}}$ its energy. This selection suppresses highly boosted, forward events that are more likely to originate from $q\bar{q}$ initial states, thereby increasing the relative contribution from gg fusion and improving sensitivity to the targeted production mechanism. Various combinations of $m_{t\bar{t}}$ and $\beta_{t\bar{t}}$ requirements were investigated, and the chosen thresholds were found to maximise the expected significance of entanglement (Appendix B).

Together, the selections described above define the *signal region* of the analysis at the detector level.

Table 7.6: Fraction of events with correctly matched, incorrectly matched or non-reconstructable down-type jets, for cases with and without a lower bound on the `spanet_hadtop_detection` score.

	correctly matched	incorrectly matched	non-reconstructable
inclusive	54.1 %	14.1 %	31.8 %
`spanet_hadtop_detection` > 0.7	73.1 %	13.5 %	13.4 %

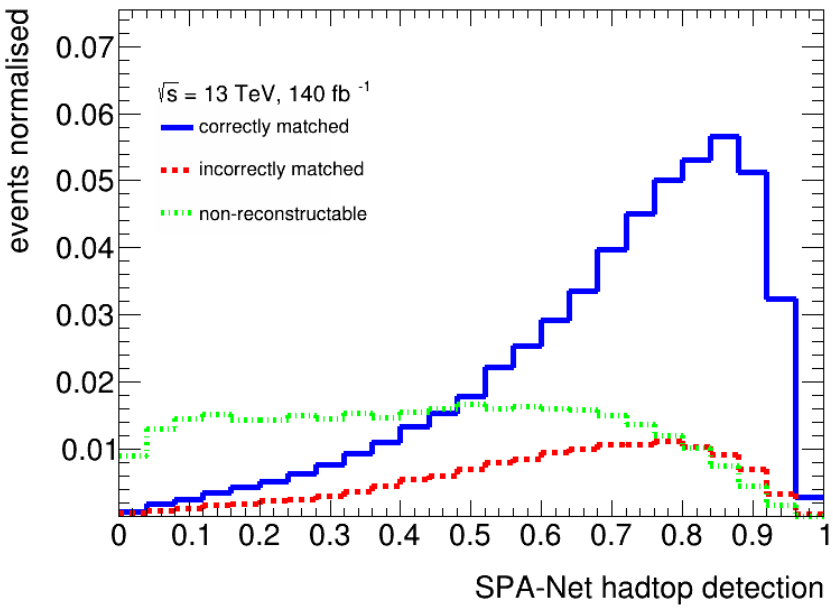


Figure 7.2: Fraction of events with correctly matched, incorrectly matched, and non-reconstructable down-type jets as a function of the `spanet_hadtop_detection` score. The fraction of correctly matched down-type jets clearly increases with higher values of this score [236].

7.6.2 Particle level

At the particle level, the reconstruction of the $t\bar{t}$ system is performed using the χ^2 minimisation technique described in Section 7.5. Once the optimal jet-parton assignment is determined, events are retained only if one of the two light jets associated with the hadronic W-boson decay contains a ghost-associated C-hadron. This requirement enhances the likelihood that the other light jet corresponds to the down-type quark. Similar to the detector-level approach, this procedure primarily targets $W \to cs$ and $W \to cd$ decays. Finally, the same selection criteria on $m_{t\bar{t}}$ and $\beta_{t\bar{t}}$ used at the detector level are applied.

7.6.3 Parton level

At the parton level, top quarks are identified as the last particles with PDG ID ± 6[1]. This definition selects the top quarks after any QCD or QED radiation and just prior to their decay. By construction, the associated W boson and b quark are defined as the direct decay products of the identified top quark. The decay mode of each top quark is determined by examining the PDG IDs of the W-boson decay products, and only events in which one top quark decays semi-leptonically and the other hadronically are retained. In order to maintain a clean and well-defined charged-lepton signature, events in which the semi-leptonic W boson decays to a τ lepton are excluded, since the subsequent τ decay introduces additional neutrinos that complicate the reconstruction of the top-quark kinematics. The parton-level event selection can therefore be summarised as follows:

- two b quarks originating directly from the top-quark decays,
- one charged lepton (electron or muon) originating from the leptonic W-boson decay,
- one neutrino originating from the leptonic W-boson decay, and
- two quarks from the hadronic W-boson decay.

To maintain consistency with the detector-level requirements, the signal region is required to have one charm quark originating from the hadronic W-boson decay. The same kinematic requirements on $m_{t\bar{t}}$ and $\beta_{t\bar{t}}$, as used at the detector and particle levels, are applied.

The selected events in the signal regions at each level (parton, particle, and detector) are used to measure the entanglement observable D through a calibration curve technique (Section 7.7). To validate the measurement strategy across different regions of the $m_{t\bar{t}}$ phase space, the measurement is additionally performed in two *validation regions* defined by $390 < m_{t\bar{t}} < 500\,\text{GeV}$ and $m_{t\bar{t}} > 500\,\text{GeV}$, while all other selection criteria remain identical to those of the signal region.

7.7 Calibration curve approach

As discussed in Section 3.2.3, the $t\bar{t}$ system is considered entangled if the measured observable $D < -1/3$. However, this criterion applies strictly at the parton level. Since experimental measurements are performed at the detector level, a method is required to infer the corresponding parton-level value

[1] PDG IDs (Particle Data Group Identifiers) are a standard, unique numbering scheme used to represent particles, their decay modes, and properties in particle physics [31]. They provide a consistent particle definition across different MC generators and are directly accessible at the parton level, since this level corresponds to a low-level event description.

of D. In high-energy physics, this procedure of recovering the underlying true distribution at parton level from what the detector actually measures is referred to as ***unfolding***.

Various unfolding methods and software tools are widely employed in ATLAS analyses [237–240]. These methods typically involve correcting the detector-level distribution to the parton level using a response matrix, thereby obtaining a differential cross-section from which the parameter of interest (in this case, D) can be extracted. The profile likelihood unfolding method, used for the boosted entanglement measurement in this thesis, follows this general approach (Section 8.6). For the threshold regime, however, an iterative bayesian unfolding method [237] was previously tested for the dilepton entanglement measurement and found to be ineffective. The main limitation came from the fact that the fiducial region sensitive to entanglement is too small relative to the experimental resolution. As a result of this, the unfolding procedure introduced a significant bias in the extracted value of D. To address this issue, a more robust alternative, known as the ***calibration curve method***, was eventually chosen for the measurement in the dilepton channel to map detector-level observables directly to their corresponding particle- or parton-level values. The same calibration curve strategy as used in the dilepton measurement is adopted in this analysis.

A calibration curve provides a method for relating definitions of the observable D across two different levels of event description. For example, D^{det} and D^{part} may represent the values of D at the detector and particle levels, respectively. The calibration curve then describes how the detector-level observable varies as a function of the particle-level observable. Conceptually, this method can be understood as a simplified version of the unfolding procedure. Rather than unfolding an entire distribution, the goal is to determine a single parameter through a mapping established between two definitions of D.

To construct the calibration curve, one begins with the nominal pair $(D^{\text{det}}, D^{\text{part}})$, where D^{det} and D^{part} are extracted from the $\cos\varphi$ distributions at detector and particle levels. In addition, several alternative simulated hypotheses $(D_i^{\text{det}}, D_i^{\text{part}})$ are generated, where i denotes the number of alternative hypotheses considered. Each hypothesis defines a point in the $(D^{\text{det}}, D^{\text{part}})$ plane, and an interpolation between these points yields a continuous calibration curve. Figure 7.3 illustrates this procedure schematically. The calibration curve can then be used to map a measured detector-level value of D (indicated by an arrow) to its corresponding particle-level value, together with an associated uncertainty.

The alternative hypotheses $(D_i^{\text{det}}, D_i^{\text{part}})$ points are obtained via a reweighting procedure. Since D is directly proportional to $\cos\varphi$, a reweighting is designed such that the $\cos\varphi$ distribution remains linear while its slope changes linearly as a function of a reweighting parameter. This procedure is purely artificial and is introduced solely to generate alternative entanglement scenarios. In this procedure, each event at a given level is assigned a reweighting factor w, defined using parton-level information as follows:

$$w(m_{t\bar{t}}^{\text{parton}}, \cos\varphi_{\text{parton}}, \chi) = \frac{1 - D_{\text{parton}}(m_{t\bar{t}}^{\text{parton}})\chi\cos\varphi_{\text{parton}}}{1 - D_{\text{parton}}(m_{t\bar{t}}^{\text{parton}})\cos\varphi_{\text{parton}}}. \tag{7.3}$$

Here, $\cos\varphi_{\text{parton}}$ denotes the value of $\cos\varphi$ at parton level for the given event, while $m_{t\bar{t}}^{\text{parton}}$ is the invariant mass of the $t\bar{t}$ system at parton level. The quantity $D_{\text{parton}}(m_{t\bar{t}}^{\text{parton}})$ is derived from a two-dimensional histogram of $\cos\varphi_{\text{parton}}$ versus $m_{t\bar{t}}^{\text{parton}}$, with 1 GeV bins in $m_{t\bar{t}}^{\text{parton}}$ and 10 bins in $\cos\varphi_{\text{parton}}$. For each event, the $m_{t\bar{t}}^{\text{parton}}$ bin is determined. The mean of the $\cos\varphi_{\text{parton}}$ distribution corresponding to that bin gives $D_{\text{parton}}(m_{t\bar{t}}^{\text{parton}})$.

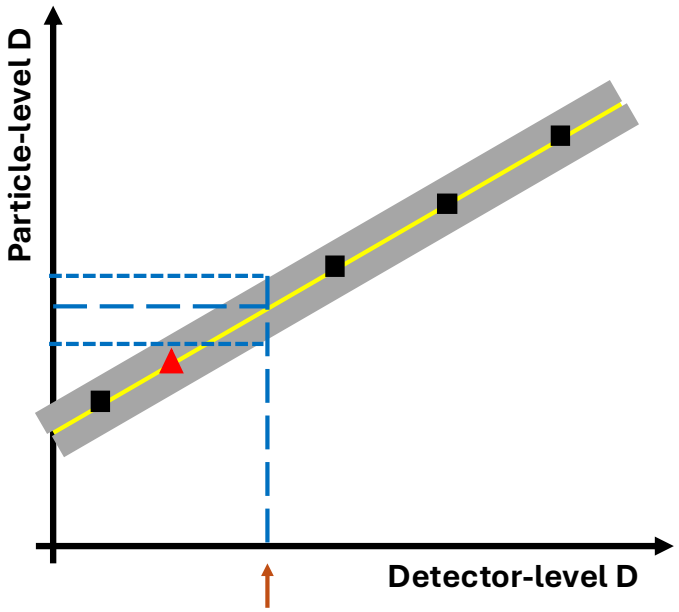


Figure 7.3: Schematic representation of the calibration curve. The nominal SM prediction obtained from simulation is shown as a red triangle, while alternative entanglement scenarios are represented by black square points derived through reweighting. These points are linearly interpolated to construct the calibration curve (yellow line). Statistical and systematic variations are evaluated at each point, resulting in additional calibration curves that define the overall uncertainty band shown in grey. A possible measured value of D is indicated by the arrow and is mapped to the corresponding particle-level value (illustrated by dashed lines) using the calibration curve. The associated uncertainty on this mapping is shown by the vertical uncertainty bands.

The parameter χ controls the scaling used to generate different entanglement scenarios. In this analysis, $\chi = \{0.4, 0.6, 0.8, 1.0, 1.2\}$, corresponding to variations of the $\cos\varphi$ slope (and therefore of D) by -60%, -40%, -20%, 0%, and $+20\%$, respectively. As an example, the reweighted distributions of the pQCD $t\bar{t}$ sample at parton level are shown in Figure 7.4.

The reweighting factors, derived using parton-level information, are subsequently applied to the corresponding events at detector and particle levels to generate four alternative hypotheses $\left(D_i^{\text{det}}, D_i^{\text{part}}\right)$, in addition to the nominal hypothesis. Together, these five points are sufficient to construct the calibration curve. The parton-level information is used consistently to define the reweighting factor such that the same transformation is applied at each level (parton, particle, or detector) entering the calibration-curve construction.

In this analysis, the calibration curve procedure is performed in two steps. First, detector-level values of D are mapped to their corresponding particle-level values through a detector-to-particle calibration curve. As a second step, the parton-level entanglement limit of $-1/3$ is propagated to the particle level using a similar calibration curve, constructed from particle- and parton-level events. The parton-level value of D is obtained from the relation $D = 3\langle\cos\varphi\rangle/\alpha_l\alpha_h$, using the NLO spin analysing powers values of 0.998 and 0.966 for the lepton and the down-type quark, repectively (Table 3.2). At the particle and detector levels, however, the down-type quark has hadronised into jets and hence is not directly observable. Therefore, the corresponding spin analysing power of the down-type jet, $\alpha_h = 0.938$, is used for the calculation of D at these levels.

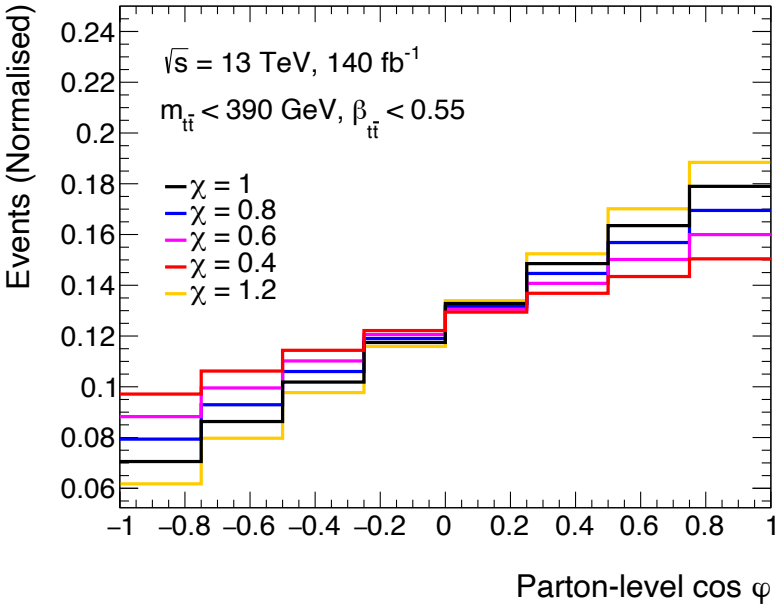


Figure 7.4: Reweighted $\cos\varphi$ distributions at the parton level, obtained using different χ values in Equation 7.3, used to construct the calibration curve.

While D is extracted here using the mean of the $\cos\varphi$ distribution, it could, in principle also be determined from the slope, according to $D = 2 \cdot \text{Slope}(\cos\varphi)/\alpha_l\alpha_h$. This alternative definition is strictly valid if the distribution is perfectly linear. In practice, however, detector effects, hadronisation of partons, and reconstruction procedures introduce deviations from linearity at the particle and detector levels. For this reason, the analysis ultimately adopts the mean-based definition of D rather than the slope-based one at all three levels.

The calibration curve is constructed using only the signal, which in this case consists of pQCD $t\bar{t}$ and $t\bar{t}$ quasi-bound-state effects. From this curve, the expected value of D at the particle level is derived. Since the signal is based on simulations, it is subject to various effects such as the modelling of signal, the definition and reconstruction of physics objects, and the algorithms used for the tagging. Uncertainties associated with background estimation are also considered, since they affect the determination of the detector-level value of D in data, as described later. These effects are collectively referred to as ***systematic uncertainties*** and account for potential biases in the prediction. In addition, ***statistical uncertainties***, arising from the finite size of the available data sample in each region, are also taken into account. The uncertainties are described in Section 7.9.

Once the calibration curve is constructed using signal simulations, the measurement proceeds on data as follows. The detector-level value of the entanglement observable D is first obtained after subtracting all background contributions from the data. The corresponding particle-level value is then determined by applying the calibration curve, translating the measured detector-level result to the corresponding particle-level value.

Although a direct detector-to-parton calibration curve could, in principle, be constructed, this two-step procedure is motivated by the presence of an uncertainty related to the PS and hadronisation

which is not well-defined, as discussed in Section 7.11. A similar behaviour of this uncertainty was also observed in the dilepton analysis [17], where the same two-step calibration curve approach was adopted. Therefore, this calibration curve strategy is also used in the present measurement.

7.8 Non-prompt and fake lepton background estimation

Backgrounds from non-prompt and fake leptons, discussed in Section 5.4, can contribute to the signal region when such leptons are misidentified as prompt ones. For example, neutral hadrons, especially through the decay $\pi^0 \to \gamma\gamma$, can generate electromagnetic showers in the calorimeters, while charged hadrons may also deposit a significant fraction of their energy electromagnetically. If a reconstructed ID track is matched to such a shower, the object can satisfy the electron identification criteria and be reconstructed as an electron candidate. The wide range of processes like these that lead to non-prompt and fake leptons, together with their strong sensitivity to detector effects, makes it difficult to model them reliably using MC simulations alone. These backgrounds are therefore estimated using data-driven techniques.

Various data-driven approaches exist to estimate this background. In this analysis, a technique known as the ***fake-factor method***, following Ref. [241], is used. As the name suggests, in this method a quantity known as the *fake factor* is derived. This factor is then applied to a set of events defined by relaxed lepton selection requirements to estimate the contribution of events in the signal region that contain fake or non-prompt leptons. The fake factor is derived in a region enriched in fake and non-prompt leptons, referred to as the *control region*. For simplicity, leptons originating from non-prompt or fake sources are collectively referred to as *fake leptons* throughout this section.

The control region is defined by requiring low missing transverse momentum, specifically events where the sum of the transverse W mass[2], m_T^W, and $E_{\mathrm{T}}^{\mathrm{miss}}$ is below 60 GeV. To maintain a phase-space which is not too different from the common object-related selection in Section 7.3, events are required to contain exactly one lepton, at least two jets, and at least two b-tagged jets satisfying the 70% WP. This relatively loose b-tagging WP is chosen to preserve a selection topology close to the signal region while increasing acceptance in the control region, thereby improving statistical precision and ensuring sufficient event yields[3] for reliable background estimation.

Events are then categorised according to the identification requirements applied to the lepton:

- **Tight events**: the lepton satisfies the tight selection criteria,
- **Loose events**: the lepton satisfies the loose but not the tight criteria,
- **Baseline events**: the sum of loose and tight events.

Leptons satisfying the loose criteria are required to pass relaxed identification requirements, with no isolation requirement applied, as summarised in Table 7.7.

Within the control region, the fraction of fake leptons satisfying the baseline selection that also satisfy the tight selection, referred to as the *fake efficiency* ϵ_{fake}, is defined as:

[2] $m_T^W = \sqrt{2p_T^l E_T^{\mathrm{miss}}(1-\cos\Delta\phi)}$, where $\Delta\phi$ is the angle between the lepton and the E_T^{miss} direction in the azimuthal plane.

[3] In this context, "yield" refers to the number of events observed or expected in a given selection region after all analysis requirements are applied.

Table 7.7: Definition of loose and tight leptons used in the non-prompt and fake lepton estimate

Lepton	Identification	Isolation
Electron (tight)	TightLH	FCTight
Electron (loose)	MediumLH	–
Muon (tight)	Medium	FCTightTrackOnly
Muon (loose)	Loose	–

$$\epsilon_{\text{fake}} = \frac{N_{\text{fake}}^{\text{tight}}}{N_{\text{fake}}^{\text{baseline}}} \tag{7.4}$$

$$= \frac{N^{\text{tight}} - N_{\text{real}}^{\text{tight}}}{N^{\text{baseline}} - N_{\text{real}}^{\text{baseline}}}. \tag{7.5}$$

Here, $N_{\text{fake}}^{\text{tight}}$ and $N_{\text{fake}}^{\text{baseline}}$ denote the number of events containing fake leptons passing the tight and baseline selections, respectively. The quantities N^{tight} and N^{baseline} denote the number of events in data passing the corresponding selections. The terms $N_{\text{real}}^{\text{tight}}$ and $N_{\text{real}}^{\text{baseline}}$ represent the expected number of events with real or prompt leptons, as estimated from simulations. Subtracting the prompt contribution from the events in data helps isolate the events with fake leptons in both selections.

The number of loose events, $N_{\text{fake}}^{\text{loose}}$, containing fake leptons satisfying the loose but not the tight citeria, is calculated as follows:

$$\begin{aligned} N_{\text{fake}}^{\text{loose}} = N^{\text{loose}} - N_{\text{real}}^{\text{loose}} &= N^{\text{baseline}} - N^{\text{tight}} - N_{\text{real}}^{\text{baseline}} + N_{\text{real}}^{\text{tight}} \\ &= N^{\text{baseline}} - N_{\text{real}}^{\text{baseline}} - (N^{\text{tight}} - N_{\text{real}}^{\text{tight}}) \\ &= N_{\text{fake}}^{\text{baseline}} - N_{\text{fake}}^{\text{tight}} \\ &= N_{\text{fake}}^{\text{baseline}} - \epsilon_{\text{fake}} \cdot N_{\text{fake}}^{\text{baseline}} \\ &= N_{\text{fake}}^{\text{baseline}}(1 - \epsilon_{\text{fake}}). \end{aligned}$$

The above expression together with Equation 7.4, results in the following:

$$\begin{aligned} N_{\text{fake}}^{\text{tight}} &= \underbrace{\frac{\epsilon_{\text{fake}}}{1 - \epsilon_{\text{fake}}}}_{\text{Fake factor F}} \cdot N_{\text{fake}}^{\text{loose}} \\ &= F \cdot N_{\text{fake}}^{\text{loose}} \\ &= F \cdot (N^{\text{loose}} - N_{\text{real}}^{\text{loose}}). \end{aligned}$$

With the estimated ϵ_{fake}, the fake factor F can be calculated. This factor is therefore used to estimate the number of events with fake leptons passing the tight selection from the corresponding number of events passing the loose and not the tight selection. The values for N^{loose} is taken from data and for $N_{\text{real}}^{\text{loose}}$ is taken from simulation.

This factor can then be used to extrapolate to the signal region (SR) and estimate the number of

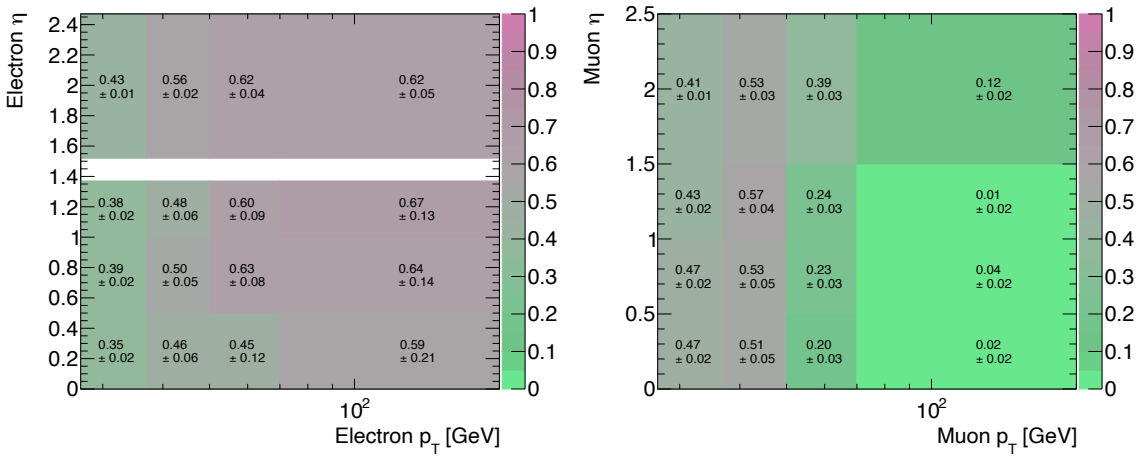


Figure 7.5: Fake efficiencies estimated in the control region for electrons (left) and muons (right) as a function of the lepton transverse momentum, p_{T}, and pseudorapidity, η.

events with fake leptons passing the tight selection, $N_{\text{fake}}^{\text{tight, SR}}$ from the corresponding number of loose events, $N_{\text{fake}}^{\text{loose, SR}}$:

$$N_{\text{fake}}^{\text{tight,SR}} = F \cdot N_{\text{fake}}^{\text{loose, SR}}$$

The fake efficiency ϵ_{fake} used to calculate the fake factor F is estimated in the control region as a function of the lepton transverse momentum p_T and pseudorapidity $|\eta|$. The resulting fake efficiencies for electrons and muons are shown in Fig. 7.5. It can be seen that the probablity for a fake lepton to satisfy the tight selection is higher for electrons than for muons at a given p_T and $|\eta|$.

This difference arises from the procedures used to reconstruct electrons and muons, as described in Sections 6.2 and 6.3, respectively. Electron reconstruction relies strongly on energy deposits left in the ECal, making electrons more susceptible to misidentification from other physics objects that produce similar electromagnetic signatures. In contrast, muon reconstruction primarily requires compatible signatures in both the ID and the MS, making it substantially more difficult for hadronic activity to mimic a muon candidate. Fake muons typically arise from in-flight decays of charged hadrons (e.g. $\pi^{\pm}$ and $K^{\pm}$), or from hadronic punch-through to the MS. These contributions are further reduced by isolation and quality requirements. As a result, the contribution from fake leptons is generally larger for electrons than for muons. The empty gap visible in the electron fake-efficiency distributions corresponds to the excluded barrel–endcap transition region of the calorimeter, introduced as part of the electron selection described in Section 7.2.1.

In practice, using parametrised fake factors $F(p_T, \eta)$, the estimate is obtained by applying event-by-event weights to the selected loose lepton sample in data, after subtraction of the prompt-lepton contribution estimated from simulation [241]:

$$N_{\text{fake}}^{\text{tight}} = \sum_{i \in \text{loose (data)}} F_i(p_T^i, |\eta^i|) - \sum_{j \in \text{loose (MC)}} w_j^{\text{MC}} F_j(p_T^j, |\eta^j|).$$

Here, F_i denotes the fake factor evaluated for lepton i as a function of its p_T and $|\eta|$. The first term corresponds to the weighted sum over loose events in data. The second term accounts for the contamination from prompt leptons, estimated using simulated events. The weights w_j^{MC} include all event weights and sample normalisation factors.

Due to statistical fluctuations and the subtraction of prompt contributions, negative yields can arise in some bins. This effect is taken into account when defining the binning of the final distributions. If the total yield in a given region is negative, the yield is set to zero (Section 7.10).

7.9 Systematic uncertainties

In this measurement, about 400 sources of systematic uncertainty are considered. These uncertainties are grouped into two main categories: *experimental uncertainties*, which are related to the detector performance and the physics objects reconstruction algorithms, and *modelling uncertainties*, which are associated with the estimation and simulation of signal and background processes. For each source of systematic uncertainty, a dedicated calibration curve is constructed. The detector-level value of D is then corrected based on this alternative calibration curve using the procedure described in Section 7.7, resulting in a shifted result, (D'_{particle}). In most cases, the systematic uncertainty is defined as the difference between the nominal result, D_{particle}, and the shifted result:

$$\Delta = D_{\text{particle}} - D'_{\text{particle}}$$

This is primarily the case for signal modelling uncertainties, where the evaluation involves a reference sample different from the nominal one (for example, using the *fast simulation* sample instead of the *full simulation*; see Section 5.2), two alternative calibration curves are constructed. The corresponding systematic uncertainty is taken as the difference between the particle-level values obtained using these two curves. For certain systematics, such as background cross-section uncertainties, the systematically shifted background sample is subtracted from data before performing the calibration. For systematics that have variations only on one side, referred to as one-sided systematics, the full difference is assigned symmetrically in both directions. In the case of asymmetric systematics, the larger variation is taken and symmetrised. All systematic uncertainties are evaluated independently at each reweighting point on the calibration curve and may therefore vary across the phase space of the D observable.

7.9.1 Experimental

Experimental uncertainties arise from several sources, including the luminosity measurement, as well as the identification efficiency, energy scale and resolution, and isolation of all relevant reconstructed physics objects used in this measurement. With the exception of the luminosity uncertainty and the signal normalisation uncertainties (associated with pQCD $t\bar{t}$ and $t\bar{t}$ quasi-bound-state effects), which only affect the overall normalisation, almost all of these uncertainties affect both the normalisation and the shape of the kinematic distributions considered in this analysis, such as $\cos\varphi$ and $m_{t\bar{t}}$. All experimental uncertainties, apart from the signal normalisation, are taken into account for all signal and background processes.

Luminosity

The uncertainty on the combined integrated luminosity of the 2015–2018 dataset is 0.83%, as described in Section 4.1.1. This uncertainty is applied as a constant normalisation shift to all MC simulated samples.

Signal normalisation uncertainties

A normalisation uncertainty of $^{+5.6\%}_{-6.1\%}$ is applied on the pQCD $t\bar{t}$ signal to reflect the uncertainty on the calculated value of the $t\bar{t}$ cross-section. For the quasi-bound-state effects, a 50% normalisation uncertainty is applied on the corresponding sample (η_t or $t\bar{t}_{\text{GFRW}}$). This is a conservative choice, motivated by the absence of an available theoretical uncertainty on the cross-section used for the normalisation of these samples (Section 5.4.6).

Pile-up

The pile-up weights introduced in Section 5.2, which are used to correct the luminosity profile of MC simulations to match that observed in data, are subject to systematic uncertainties. These uncertainties are evaluated by varying the assumed average number of interactions per bunch crossing in the pile-up weight calculation. The resulting variations define the corresponding systematic shifts in both the upward and downward directions.

Electron and muon

The dominant uncertainties associated with electrons and muons stem from the scale factors applied to account for reconstruction, identification, isolation, and trigger efficiencies, as discussed in Sections 7.2.1 and 6.3. For electrons, separate systematic uncertainties are assigned to each of these components. For muons, uncertainties are similarly applied to the identification, isolation, trigger, and track-to-vertex association efficiencies, accounting for both their statistical and systematic components.

Additional uncertainties originate from the calibration of the energy scale and resolution. For electrons, these are evaluated through variations in the energy scale and resolution. For muons, uncertainties in the momentum scale are considered, along with contributions from the *sagitta* (track curvature) measurement. Furthermore, uncertainties in the momentum scale and resolution of CB muons are also taken into account.

Jets

The different methods involved in the calibration of jets (Section 6.4.5) lead to a large number of systematic uncertainties [214]. The uncertainties considered in this analysis are summarised below.

The uncertainties related to the JES arise from the different steps used in the jet energy calibration chain: η-intercalibration, Z+jet balance, γ+jet balance, multijet balance, and single-particle calibrations at high p_T. Additional flavour-dependent JES uncertainties, including the flavour response and flavour composition components, account for differences in detector response to quark and gluon jets. Quark and gluon jets differ in their particle multiplicity and energy distribution, with gluon jets typically being wider and containing more low-energy particles.

Altogether, the resulting JES uncertainties lead to a large set of correlated components. Several schemes exist to derive a reduced set of uncertainties, to allow simplified descriptions with minimal loss of correlation information. An *eigenvector decomposition* method is first performed on the covariance matrix of these uncertainty components. This procedure effectively diagonalises the correlation structure and yields a set of orthogonal components ordered by their significance. The dominant components are retained as independent sources, while the remaining subleading contributions are combined. This combination is achieved through a *category reduction* scheme, which groups components according to their origin [214], such as statistical, modelling, detector, flavour, or mixed sources. After this reduction, the JES uncertainties are described by a total of 35 independent components.

Uncertainties related to the JER are incorporated by smearing jets with a Gaussian function of width σ_{smear}, defined as [214]:

$$\sigma^2_{\text{smear}} = (\sigma_{\text{nominal}} + |\sigma_{\text{syst}}|)^2 - \sigma^2_{\text{nominal}},$$

where σ_{nominal} denotes the nominal JER of a sample, and σ_{syst} corresponds to a one-standard-deviation variation of the considered uncertainty component. To ensure the correct propagation of these systematic uncertainties to the analysis, smearing is applied to the simulation when $\sigma_{\text{syst}} > 0$ and to the data when $\sigma_{\text{syst}} < 0$. In cases where smearing the data is not reliable due to limited statistics, pseudo-data may be smeared instead. This procedure affects only the systematic variation of the JER component while preserving the nominal JER in data. When the JER is lower in simulation than in data, this is already accounted for by the initial smearing of the simulation (Section 6.4.5). However, an additional systematic uncertainty is required to cover cases where the JER is lower in data than in simulation. This uncertainty is taken as the full difference between the JER in data and in simulation. An eigenvector decomposition, similar to that used for JES, is employed to obtain a reduced set of JER uncertainties. The final set consists of 13 independent components.

A single uncertainty is applied on the JVT discriminant, accounting for statistical uncertainties and an additional 30% uncertainty in the estimation of the residual pile-up contamination after pile-up suppression.

Flavour tagging

As described in Section 6.4.6, the mismodelling of the DL1r discriminant distribution in simulation is corrected by applying flavour-dependent scale factors derived from data, such that the b-tagging efficiency in simulation matches that observed in data. Corrections applied to simulated samples to reproduce the performance of the DL1r algorithm in data, along with their associated uncertainties, are determined separately for each jet flavour generally in bins of jet p_T and in the five PCBT score bins. This procedure is referred to as the PCBT calibration. Uncertainties on the b-, c-, and light-flavour tagging efficiency scale factors for the PCBT calibration arise from a number of sources specific to the methods used to derive these scale factors, including uncertainties related to track reconstruction and the modelling of PFLow jets. The number of components associated with these uncertainties is significantly reduced using the *eigenvector variation method*, in which the systematic variations are obtained by diagonalising the covariance matrix that encodes correlations across bins of the kinematic variables in which the scale factors are derived. This diagonalisation provides orthogonal variations whose size is given by the square root of the corresponding eigenvalue. The number of eigenvector

variations obtained for the PCBT calibration are 45, 20, and 20 for the b-, c-, and light-flavour PCBT scale factors, respectively. To account for the extrapolation of the scale factors to the high-p_T regime, an additional 4, 2, and 4 components, respectively, are included.

Missing transverse momentum

Uncertainties in the missing transverse momentum (Section 6.5), $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$, originate from the various inputs entering its definition. A fraction of these uncertainties is propagated from the calibration of the reconstructed physics objects contributing to the hard-event term, such as small-R jets, electrons, and muons. The remaining contribution arises from uncertainties associated with the soft term. The uncertainty on $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ is therefore propagated in the fit using three components. One component corresponds to the scale uncertainty of the soft term, while the other two describe its resolution, defined along the directions parallel and perpendicular to the transverse momentum of the hard term [224].

7.9.2 pQCD $t\bar{t}$ modelling systematics

The nominal modelling of the pQCD $t\bar{t}$ sample is based on the Powheg Box+ Pythia generator setup described in Section 5.4.1, which carries a set of theoretical and modelling uncertainties. These arise from choices in the ME calculation, PS modelling, and related parameters. To evaluate the impact of these assumptions and account for the associated systematic uncertainties, two approaches are considered: *alternative generators* and *weights*. For uncertainties estimated using alternative generators or samples produced with alternative Powheg Box+ Pythia generator settings, the systematic effect is obtained by comparing the values of D extracted using the calibration curve constructed from the nominal sample with those obtained using the calibration curve derived from the corresponding alternative sample. For the systematic uncertainties that use weights, variations are applied to the nominal sample. The systematic uncertainty is then determined from the difference between the value of D obtained using the nominal calibration curve and that obtained after applying the weights.

An extensive range of uncertainties related to the MC modelling of the pQCD $t\bar{t}$ signal process are taken into account, as summarised below. Unless otherwise stated, these uncertainties affect both the shape and the normalisation of the pQCD $t\bar{t}$ differential distributions.

NNLO-QCD scale variation

The $t\bar{t}$ reweighting procedure described in Section 5.4.1 introduces additional uncertainties affecting the reweighted prediction. These are split into two components: one related to the NNLO-QCD reweighting, and the other to the NLO-EW part. The effect of the NNLO-QCD reweighting is estimated by reweighting the nominal $t\bar{t}$ pQCD sample to alternative NNLO-QCD+NLO-EW predictions obtained by varying μ_R and μ_F by factors of 2 and 0.5. The systematic uncertainty is taken as the envelope of all resulting variations. All μ_R and μ_F combinations are considered except $\mu_R = 2, \mu_F = 0.5$ and vice versa [242]. The envelope is dominated by the variation of μ_R.

NLO-EW correction scheme

The second component of the reweighting uncertainty originates from the NLO-EW corrections. It is evaluated by comparing two schemes for combining NLO-EW and NNLO-QCD corrections: additive

and multiplicative [243]. In the multiplicative scheme, which is the nominal scheme, the combined correction is of the form $(1 + \delta_{\text{QCD}})(1 + \delta_{\text{EW}})$, while in the additive scheme it is $(1 + \delta_{\text{QCD}} + \delta_{\text{EW}})$, where δ_{QCD} and δ_{EW} denote the relative QCD and EW corrections, respectively. The systematic uncertainty is defined as the difference between these two schemes.

Top-quark mass

The systematic uncertainty associated with the top-quark mass, m_{top}, is estimated using two alternative Powheg+Pythia8 samples in which m_{top} is varied by ±0.5 GeV. This variation corresponds to the total top-quark mass uncertainty obtained in the ATLAS combination [244].

Lineshape

The lineshape uncertainty is associated with the modelling of top-quark decays and off-shell effects. It arises from differences in the treatment of spin correlations in $t\bar{t}$ production across generators. In principle, this uncertainty is assessed by comparing the sum of the $t\bar{t}$ sample and the tW sample to an alternative sample generated with the bb4l setup [245]. This alternative bb4l model includes off-shell and non-resonant contributions as well as exact spin correlations at NLO accuracy and simulates the inclusive production of $b\bar{b}ll^{+}\nu\bar{\nu}$ final states. In this way, this model takes into account the interference effects between $t\bar{t}$ and tW. In this comparison, all samples use the Pythia setting in which the top quark participates in the recoil after the first gluon emission.

However, for the evaluation of this uncertainty in the lepton+jets channel, the corresponding bb4l sample is not available. Therefore, the direct comparison described above is not possible. Instead, the nominal Powheg+Pythia sample is reweighted using weights derived from the comparison between the corresponding samples and bb4l in the dilepton channel. The resulting reweighted sample is then compared to the nominal Powheg+Pythia prediction to estimate this uncertainty, which provides the best available approximation given the current availability of the bb4l simulation.

PDFs

The systematic uncertainty associated with the choice of PDF is evaluated using the PDF4LHC 30 eigenvector set [65]. For each eigenvector, the difference between the nominal PDF weight and the variation is calculated and then symmetrised. These eigenvector variations are available as generator weights in the nominal Powheg sample.

PS and hadronisation

This uncertainty is related to the choice of the PS and hadronisation scheme and is estimated by comparing the nominal prediction with a prediction obtained using an alternative pQCD $t\bar{t}$ sample which is also produced using the Powheg Box v2 generator but interfaced with Herwig 7.2 [131, 132]. The default set of tuned parameters [132], the MMHT2014$_{\text{LO}}$ PDF set [246], and the same Powheg setting as the nominal sample are used.

ME/PS matching

The ME/PS matching uncertainty is evaluated by comparing the nominal Powheg+Pythia sample to an alternative sample in which the $p_{\mathrm{T}}^{\mathrm{hard}}$ setting in Pythia is varied. This parameter controls the matching between the ME calculations of the hard scatter and the PS. The default value, $p_{\mathrm{T}}^{\mathrm{hard}} = 0$, corresponds to Powheg providing the p_T of its hardest emission to Pythia to define the matching scale. While this choice is unambiguous for events without additional real emissions, alternative definitions are possible when such emissions are present. In particular, the scale can instead be defined using the p_T of the emission generated by Powheg, corresponding to $p_{\mathrm{T}}^{\mathrm{hard}} = 1$. The matching uncertainty is therefore estimated by comparing the nominal sample with $p_{\mathrm{T}}^{\mathrm{hard}} = 0$ to the alternative sample with $p_{\mathrm{T}}^{\mathrm{hard}} = 1$, following the prescription described in [247].

ISR h_{damp}

This uncertainty is associated with the ME/PS matching scale, h_{damp}, introduced in Section 5.4, which controls the amount of ISR handled by the ME and the PS. The uncertainty is estimated by comparing the nominal prediction with a Powheg+Pythia8 sample in which the value of h_{damp} is doubled.

ISR

The uncertainty on the modelling of ISR is primarily driven by variations in the strong coupling constant, α_s. These effects are accounted for by using event weights provided in the A14 tune of the nominal Powheg sample. Specifically, the weights `Var3cUp` and `Var3cDown` are applied [139], as they were shown to encompass the data in ISR-sensitive observables within the tune.

FSR

The uncertainty associated with the modelling of FSR is evaluated by varying the renormalisation scale used in the FSR part of the PS. The scale is varied up and down by a factor of two, resulting in two variations that define the FSR uncertainty component.

Recoil-to-top

The b-quarks produced in the top-quark decay, $t \to Wb$, can radiate gluons through QCD radiation. These radiations are modelled in the simulation of the decay. The first emission, generated by the ME calculation in Powheg, has a well-defined kinematic configuration. However, subsequent emissions are produced by the PS in Pythia, where momentum conservation requires each emission to be balanced by a recoil. No unique choice exists for the recoiling particle. Therefore, an ambiguity arises in the choice of the recoil partner for gluon emissions beyond the first emission [248]. This choice can affect out-of-cone radiation associated with reconstructed b-tagged jets, and can therefore impact the reconstructed top-quark mass. In the nominal Pythia setup, the recoil is assigned to the nearest coloured parton, in this case the b-quark. An uncertainty for this is assessed by using an alternative sample with Pythia settings in which the top quark participates in the recoil, hence the name "recoil-to-top".

Colour reconnection

Colour reconnection denotes the strong interaction between colour-singlet parton systems originating at different stages of the event evolution and is a key feature of the UE. As briefly described in Section 5.2, in MC event generators such as Pythia, quarks are connected to other partons via colour strings, while gluons are each connected to two other partons, since they carry both a colour and an anti-colour charge. The presence of MPIs introduces many additional coloured partons in the final state, resulting in a dense network of colour strings. This influences the development of the PS and the subsequent hadronisation. To resolve the ambiguities arising from overlapping colour connections between different MPIs and the hard-scattering process, colour reconnection models are introduced to reassign colour connections in a physically motivated way. To assess the associated uncertainty, three different models and their corresponding tunes are considered, denoted as CR1, CR2, and CR0 [249]. These tunes are derived from fits to UE data and differ in both the modelling approach and parameter settings, including parameters related to MPI. The CR1 model is a QCD-based model that enhances the formation of dipoles containing three quarks, leading to increased baryon production. The CR2 model, known as the gluon-move model, restricts reconnections to gluons, which are rearranged to reduce the overall string length. The CR0 model is the nominal MPI model with a tune matching that of CR1 and CR2. The predictions obtained using the CR1 and CR2 models are compared to those from the CR0 model, and the largest deviation is taken as the systematic uncertainty. For the entanglement measurement, the CR2 model yields the largest variation. The comparison is performed with respect to CR0 rather than the nominal sample to avoid including additional uncertainties associated with the tuning procedure.

Underlying Event

The difference in UE modelling is assessed by comparing the nominal sample with two samples that use the same algorithms for ME and PS as the nominal sample, but different choices of the Var 1 parameter of the A14 set of tuned parameters [139].

7.9.3 Modelling systematics of $t\bar{t}$ quasi-bound-state effects

Modelling uncertainties are also considered for $t\bar{t}$ quasi-bound state production. These include uncertainties related to the modelling of the PS and hadronisation, as well as ISR and FSR. They are evaluated by comparing the nominal predictions to those obtained from samples generated with alternative settings, following the same procedure as for the corresponding uncertainties in the $t\bar{t}$ pQCD sample. These uncertainties are treated as uncorrelated with those in the pQCD sample, given the different orders in α_s at which the processes are generated.

These modelling uncertainties are applied only to the simplified model, η_t. As already mentioned before, the η_t model provides a first approximate description of quasi-bound-state effects and is not intended as a final or complete theoretical description. Therefore, only a limited set of systematic uncertainties is considered for this model at this stage.

No modelling uncertainties are assigned to the nominal model, $t\bar{t}_{\text{GFRW}}$. At the time of writing, the corresponding systematic variations were not yet available for the lepton+jets final state and are therefore not included in the fit. These will be incorporated in future iterations where models with improved modelling are expected to become available. Nevertheless, it has been verified that the impact of the reduced set of modelling uncertainties in the simplified model is negligible in the signal

region used in this analysis (Section 7.11), which covers a broader $m_{t\bar{t}}$ range compared to the dedicated threshold analyses [21, 22]. Therefore, the missing modelling uncertainties are not expected to have a significant impact on the final measurement.

7.9.4 Background systematics

Associated production of $t\bar{t}$ with heavy flavour ($t\bar{t}+b$ and $t\bar{t}+c$)

In $t\bar{t}$ events, additional heavy-flavour quarks (b or c) can be produced. Events in which a $t\bar{t}$ pair is produced in association with an extra b- or c-quark, referred to as $t\bar{t}+b$ and $t\bar{t}+c$, respectively, are identified within the nominal pQCD $t\bar{t}$ sample using particle-level information. The relative contribution of the $t\bar{t}+b$ ($t\bar{t}+c$) subsample is varied by 20% (40%) to account for uncertainties in the predicted cross-sections of these processes [250, 251].

Single-top production

The single-top background in this analysis is dominated by tW production. The main modelling uncertainty in this component arises from the treatment of overlap between tW and $t\bar{t}$ production, quantified by comparing the DR and DS schemes discussed in Section 5.4.2. This uncertainty is evaluated by comparing the nominal tW sample generated with the DR scheme to an alternative sample generated with the DS scheme. Additional uncertainties related to the variations of the $p_{\mathrm{T}}^{\mathrm{hard}}$ and h_{damp} parameters, as well as the PS model, are estimated following the same procedure as for $t\bar{t}$ production. For the PS uncertainty, a tW sample using Herwig 7.2 for the shower and hadronisation is employed. A normalisation uncertainty of 6% is applied to account for the theoretical uncertainty on the tW cross-section. The same normalisation uncertainty is also applied to the minor single-top backgrounds from s- and t-channel production modes.

Uncertainties related to other backgrounds

A conservative 50% normalisation uncertainty is applied to the contributions from $t\bar{t}H$ and $t\bar{t}V$ to account for potential mismodelling effects. For the W+jets background, a 50% normalisation uncertainty is assigned to cover possible mismodellings of both the inclusive cross-section and the flavour composition. In the case of the Z+jets background, a conservative 50% normalisation uncertainty is applied. This accounts for the ±5% uncertainty on the inclusive cross-section as well as additional acceptance uncertainties, in particular those related to the modelling of events with high jet multiplicities. A similar 50% normalisation uncertainty is assigned to the diboson background to account for potential mismodelling in the production of additional jets [252] and heavy-flavour jets [253]. Finally, a conservative 50% normalisation uncertainty is applied to the data-driven estimate of the fake or non-prompt lepton background.

7.10 Distributions and event yield

The detector-level $m_{t\bar{t}}$ distribution, reconstructed via SPANet using the procedure described in Section 7.4.2, is shown in Figure 7.6. The uncertainty bands shown include statistical contributions as well as all systematic uncertainties discussed in Section 7.9. Individual background components are also displayed. Overall agreement between data and the prediction is observed, and is well within the evaluated uncertainties. This comparison between data and simulation provides a validation of both the modelling of the selected events and the performance of the $t\bar{t}$ kinematic reconstruction. Selections on this reconstructed $m_{t\bar{t}}$, together with the requirements $\beta_{t\bar{t}} < 0.55$ and `spanet_hadtop_detection` > 0.7, are used to define the signal and validation regions employed in this measurement.

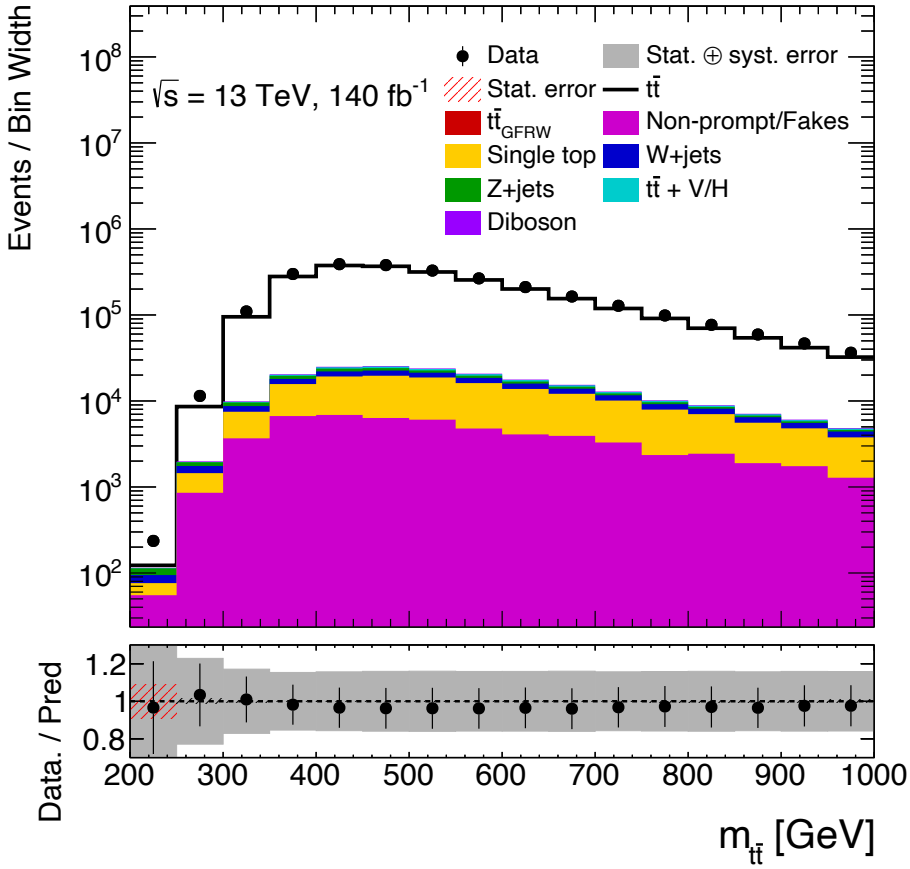


Figure 7.6: Distribution of the invariant mass of the $t\bar{t}$ system, $m_{t\bar{t}}$, at the detector level, showing the breakdown of signal and background contributions. The ratio between data and prediction is shown in the lower panel. The uncertainties include statistical (dashed red) and all systematic contributions considered in this analysis. No selections on $\beta_{t\bar{t}}$ and SPANet performance variables are applied. The y-axis is in the logarithmic scale.

The detector-level $\cos\varphi$ distribution, which is central to the extraction of D, is shown in Figure 7.9 for the signal region. The corresponding transverse momentum, p_T, distributions of the lepton and the down-type jet, which are central to defining this angle are shown in the laboratory frame for the signal region in Figure 7.8. Good agreement between data and prediction is observed for these distributions,

providing further confidence in the modelling of the kinematic quantities entering the entanglement measurement.

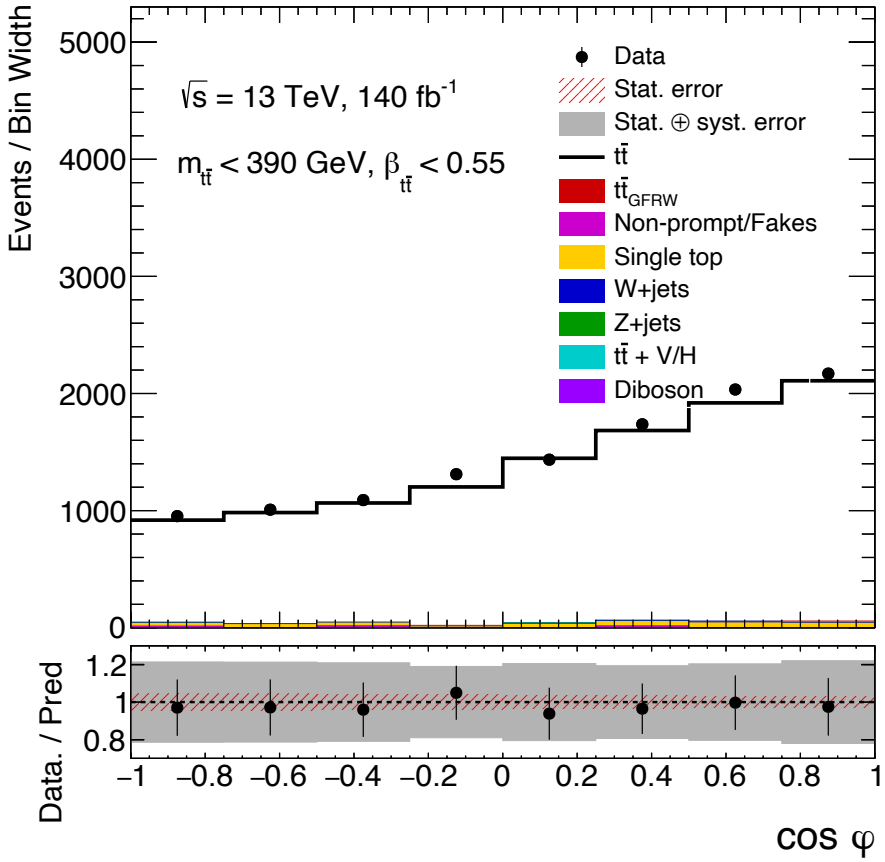
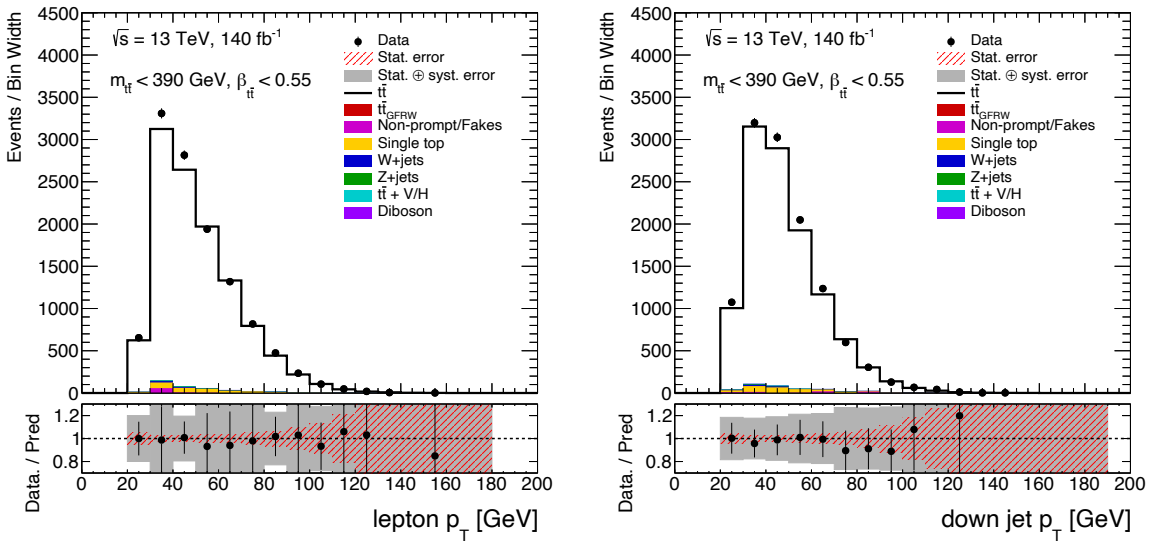


Figure 7.7: Distribution of the angular variable $\cos\varphi$ at the detector level in the signal region ($m_{t\bar{t}} < 390\,\text{GeV}$), showing the breakdown of signal and background contributions. The ratio between data and prediction is shown in the lower panel. The uncertainties include statistical (dashed red) and all systematic contributions considered in this analysis.

The corresponding $\cos\varphi$ distributions in the validation regions are shown in Figure 7.9. The $t\bar{t}$ signal is observed to dominate across all regions, contributing more than 90% of the total yield in each of the three regions. This is further quantified in Table 7.8, which summarises the event yields. Their sum defines the total expected yield reported in the table. The corresponding yields observed in data are also included, together with the data-to-expectation ratios. The level of agreement is strongest in the low $m_{t\bar{t}}$ signal region, with good agreement also found in both validation regions. As expected, $t\bar{t}_{\text{GFRW}}$ has its largest contribution in the signal region.

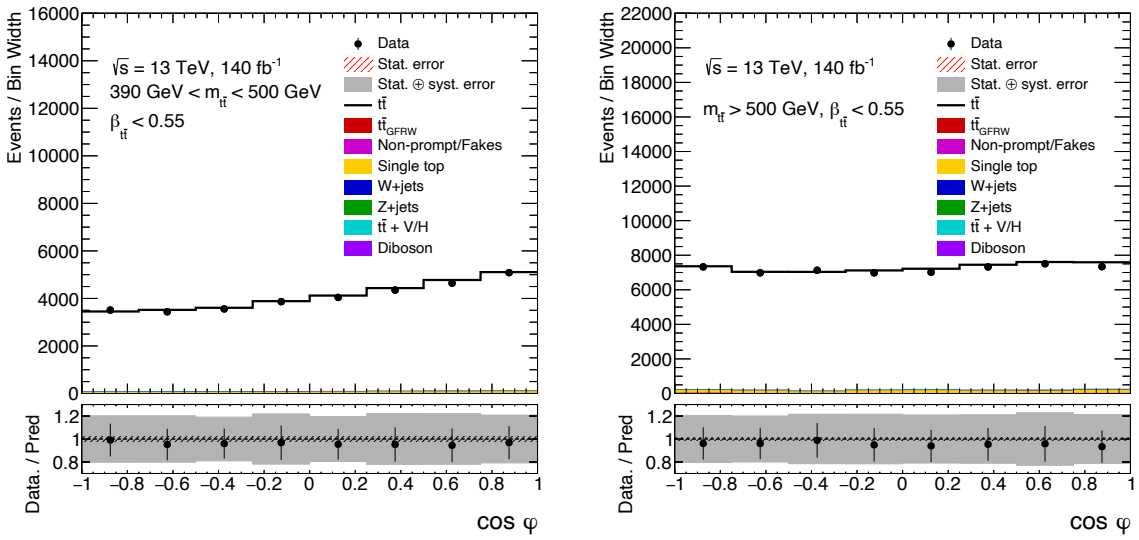
It is worth noting that the event yield for non-prompt and fake backgrounds in one of the validation regions is reported as zero. This value was set explicitly, as the data-driven estimation procedure (Section 7.8) resulted in a negative yield in this region, which is unphysical in the context of this analysis. The yield was therefore set to zero under the assumption that this background contribution is negligible in this region.



(a) Lepton

(b) Down-type jet

Figure 7.8: Detector-level distributions of the transverse momentum, p_T, of the lepton and the down-type jet, which serve as spin analysers used in the construction of the angular observable $\cos\varphi$, in the signal region. The ratio between data and prediction is shown in the lower panel. The uncertainties include statistical (dashed red) and all systematic contributions considered in this analysis



(a) Validation region (390 GeV $< m_{t\bar{t}} <$ 500 GeV)

(b) Validation region ($m_{t\bar{t}} >$ 500 GeV)

Figure 7.9: Distributions of the angular variable $\cos\varphi$ at the detector level in the validation regions, showing the breakdown of signal and background contributions. The ratio between data and prediction is shown in the lower panel. The uncertainties include statistical (dashed red) and all systematic contributions considered in this analysis.

Table 7.8: Event yield summary for the signal and validation regions. The quoted uncertainties on the predicted event yields include both statistical and systematic contributions, summed in quadrature.

Process	$m_{t\bar{t}} < 390\,\text{GeV}$	$390\,\text{GeV} < m_{t\bar{t}} < 500\,\text{GeV}$	$m_{t\bar{t}} > 500\,\text{GeV}$
$t\bar{t}$	11300 ± 600	32900 ± 1800	58400 ± 3200
$t\bar{t}_{\text{GFRW}}$	270 ± 60	81 ± 16	19.1 ± 3.5
Single-top-quark	249 ± 18	654 ± 33	1390 ± 60
W+jets	41 ± 32	65 ± 25	91 ± 30
Z+jets	21 ± 9	34 ± 11	39 ± 10
$t\bar{t}+V/H$	25 ± 5	90 ± 16	230 ± 40
Diboson	3.4 ± 1.6	7.0 ± 3.3	9 ± 4
Non-prompt/Fakes	60 ± 50	0.0 ± 0.0	100 ± 90
Total expectation	12001 ± 600	33800 ± 1800	60300 ± 3200
Data	11738	32492	57605
Data/Expectation	0.98	0.96	0.95

7.11 Results

The entanglement observable D is extracted in the signal region as well as in the two validation regions and, using the calibration curve technique described in Section 7.7, mapped to its particle-level value. All relevant systematic uncertainties are included in each region. The result obtained from data is referred to as the *observed* value, while the prediction derived from simulation is called the *expected* value.

7.11.1 Calibration curves and extraction of D using the nominal signal model

The calibration curve is constructed in the signal and validation regions using the nominal signal, consisting of the pQCD $t\bar{t}$ sample and the quasi-bound-state contribution $t\bar{t}_{\text{GFRW}}$. The resulting particle-level values of the entanglement observable D obtained using this curve are:

$$D_{\text{observed}} = -0.357 \pm 0.008[\text{stat.}]^{+0.033}_{-0.040}[\text{syst.}],$$
$$D_{\text{expected}} = -0.365 \pm 0.008[\text{stat.}]^{+0.033}_{-0.040}[\text{syst.}],$$

in the signal region of $m_{t\bar{t}} < 390\,\text{GeV}$,

$$D_{\text{observed}} = -0.201 \pm 0.004[\text{stat.}]^{+0.031}_{-0.028}[\text{syst.}],$$
$$D_{\text{expected}} = -0.218 \pm 0.004[\text{stat.}]^{+0.031}_{-0.028}[\text{syst.}],$$

in the validation region of $390 < m_{t\bar{t}} < 500\,\text{GeV}$, and

$$D_{\text{observed}} = -0.128 \pm 0.003[\text{stat.}] \pm 0.019[\text{syst.}],$$
$$D_{\text{expected}} = -0.146 \pm 0.003[\text{stat.}] \pm 0.019[\text{syst.}],$$

in the validation region of $m_{t\bar{t}} > 500\,\text{GeV}$. Here, D_{observed} and D_{expected} denote the observed and expected values of D at particle level.

The calibration curve for the signal region is shown in Figure 7.10. The black marker represents the value of D extracted from the data, while the square markers correspond to the alternative entanglement scenarios used to construct the curve (Section 7.7). The red marker denotes the SM expectation of D as obtained from the nominal signal. The yellow and grey bands represent the statistical and total (statistical plus systematic) uncertainties, respectively. The corresponding calibration curves for the validation regions are shown in Figure 7.11.

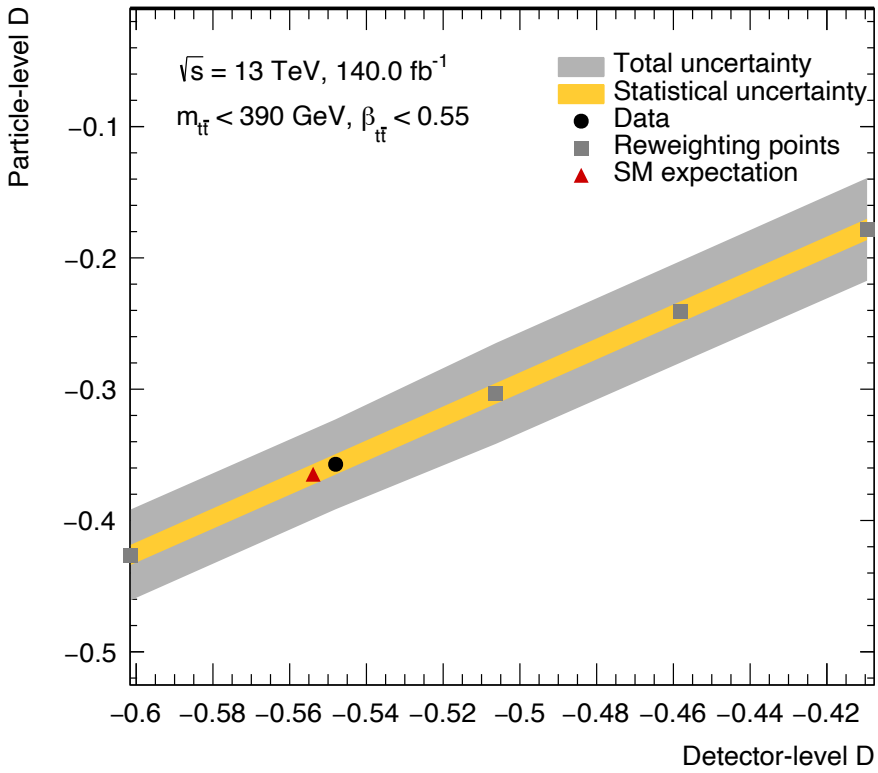
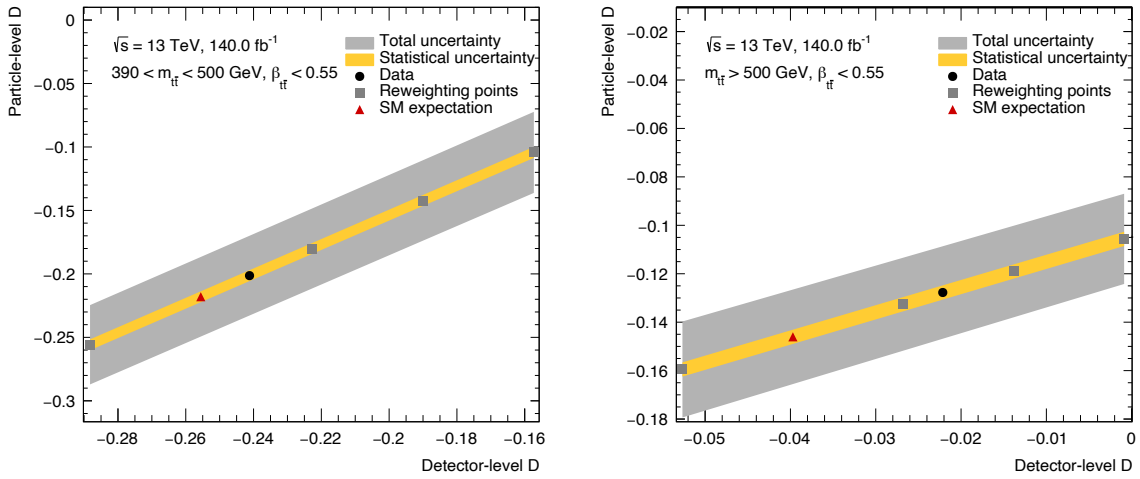


Figure 7.10: Calibration curve in the signal region ($m_{t\bar{t}} < 390\,\text{GeV}$). The black filled circular marker represents the data while the SM expectation is shown by the red triangle marker. The grey square markers represent the alternative entanglement scenarios introduced in Section 7.7.

Table 7.9 summarises the impact of the systematic uncertainties discussed in Section 7.9, grouped by category and summed in quadrature, together with the statistical uncertainty, on the observed and expected values of D in the signal region. It can be seen that the pQCD $t\bar{t}$ modelling has the largest impact on both the observed and expected values. A more detailed breakdown of the ten dominant individual sources of systematic uncertainty affecting the observed value of D in the signal region is presented in Table 7.10. The leading contributions arise from uncertainties related to the ME/PS matching, FSR, PS and hadronisation, as well as colour reconnection. In addition, several components



(a) Validation region ($390 < m_{t\bar{t}} < 500\,\text{GeV}$)

(b) Validation region ($m_{t\bar{t}} > 500\,\text{GeV}$)

Figure 7.11: Calibration curves in the validation regions, similar to Figure 7.10. Note that the axes scales differ.

of the JER uncertainty contribute significantly. Furthermore, one of the eigenvector variations of the light-flavour PCBT scale factors is among the dominant sources. This impact could originate from the strong sensitivity of the analysis to the correct identification of the down-type jet originating from the hadronic W-boson decay.

To assess whether the observed (expected) value is below the entanglement limit, the parton-level limit of $-1/3$ is translated to the particle level using a parton-to-particle calibration curve in the signal region, as shown in Figure 7.12. The point corresponding to this limit is highlighted with a star. The associated statistical uncertainty is negligible, while the systematic uncertainties are significantly smaller than those of the detector-to-particle calibration curve.

In the construction of this curve, only signal modelling uncertainties are considered, as experimental uncertainties do not affect particle- or parton-level observables. The PS and hadronisation uncertainty is not included as one of the systematic uncertainties in this case. This is because the comparison between the nominal Powheg+Pythia setup and the alternative Powheg+Herwig configuration, used to estimate this uncertainty, does not isolate a single modelling effect. In addition to differences in the PS, the comparison also includes differences in the hadronisation model and several other generator settings that are not directly comparable. In particular, Pythia employs the Lund string model with a p_{T}-ordered shower, while Herwig uses a cluster hadronisation model with angular-ordered showering (Section 5.2). Therefore, the resulting uncertainty is not well-defined, although this comparison currently provides the best available uncertainty estimate of the PS modelling.

A comparison between the calibration curves obtained using the nominal Powheg+Pythia setup and the alternative Powheg+Herwig configuration, is shown in Figure 7.13(a). The difference between the two curves is significantly larger than that observed for the corresponding detector-to-particle calibration curve (Figure 7.13(b)), a feature also reported in the dilepton entanglement analysis [17]. These differences affect the $\cos\varphi$ distributions at the particle level, while their impact at the parton level remains negligible.

Table 7.9: A summary of the grouped sources of systematic uncertainty in terms of their absolute and relative impact on the observed and expected values of D after calibration from detector to particle level.

	$D_{\text{observed}} = -0.357$		$D_{\text{expected}} = -0.365$	
Systematic source	Absolute Size	Relative Size [%]	Absolute Size	Relative Size [%]
pQCD $t\bar{t}$ modelling	0.0234	6.56	0.0234	6.43
Electron	0.0012	0.34	0.0012	0.34
Muon	0.0005	0.14	0.0005	0.14
Jets	0.0212	5.95	0.0212	5.82
Flavour tagging	0.0069	1.93	0.0071	1.94
Pile-up	0.0015	0.42	0.0015	0.41
$E_{\text{T}}^{\text{miss}}$	0.0025	0.71	0.0025	0.69
Backgrounds	0.0071	2.0	0.0071	1.96
$t\bar{t}_{\text{GFRW}}$ cross-section	0.0017	0.47	0.0017	0.45
pQCD $t\bar{t}$ cross-section	0.0002	0.06	0.0002	0.06
Stat.	0.0078	2.18	0.0078	2.14
Syst.	0.0334	9.34	0.0334	9.15
Total	0.0343	9.59	0.0343	9.4

Table 7.10: Ranking of the ten largest sources of systematic uncertainty in terms of their absolute and relative impact on the observed value of D after calibration from detector to particle level. NP-(X) denotes the X-th component within the set of jet or flavour-tagging uncertainties, as introduced in Section 7.9.1. The Jet JER Data vs. MC uncertainty corresponds to the component associated with differences in JER modelling between data and simulation.

Name of the Uncertainty	Absolute Size	Relative Size [%]
ME/PS matching (p_T^{hard})	0.0178	4.97
FSR	0.0107	2.99
Jet pile-up ρ topology	0.0098	2.75
Jet JER NP-1	0.0094	2.65
Parton Shower	0.0076	2.14
Jet JER NP-2	0.0063	1.75
Jet JER Data vs. MC	0.0060	1.67
Colour reconnection	0.0058	1.62
b-tag (Light) NP-0	0.0053	1.49
Jet JER NP-6	0.0051	1.43

For this reason, the entanglement limit is evaluated separately for the Powheg+Pythia and Powheg+Herwig configurations. The Powheg+Pythia entanglement limit is obtained to be $-0.229^{+0.008}_{-0.011}$[syst.], with a negligible statistical uncertainty of $4.3 \cdot 10^{-4}$. The quoted systematic

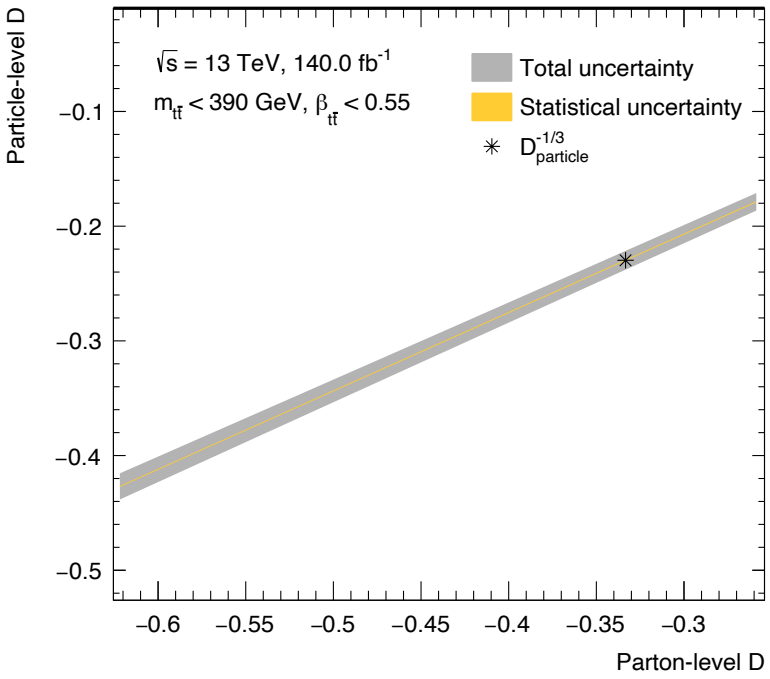


Figure 7.12: Calibration curve mapping the parton-level D value to the particle-level value, including statistical and modelling systematic uncertainties (except the PS and hadronisation uncertainty). The entanglement limit $D = -1/3$ propagated from parton to particle level, denoted by $D^{-1/3}_{\text{particle}}$, is marked with a star.

uncertainty includes all modelling uncertainties except those related to PS and hadronisation. A summary of the most significant systematic uncertainties affecting this value is provided in Table 7.11. The corresponding Powheg+Herwig entanglement limit is −0.193. No uncertainty is assigned to this value, as it is used only as an alternative modelling scenario for comparison purposes.

The observed and expected values of D lie well below both entanglement limits in the signal region. In contrast, the values obtained in the validation regions remain above the corresponding thresholds, as expected, since these regions are not expected to be enriched in entangled states. A more detailed discussion is provided in Section 7.11.3.

7.11.2 Results without quasi-bound-state effects

As in the dilepton analysis [17], the calibration curve is constructed using only the pQCD $t\bar{t}$ sample as signal, without including $t\bar{t}$ quasi-bound-state effects in the prediction. The resulting particle-level value of D allows a direct comparison with the corresponding dilepton results, which did not include any $t\bar{t}$ quasi-bound-state effects. Furthermore, it provides a straightforward assessment of the impact of including $t\bar{t}$ quasi-bound-state effects on the measured value of D. The resulting particle-level

Table 7.11: Ranking of the impact of the ten most dominant sources of systematic uncertainty on the calibration of the entanglement limit $D = -1/3$ from the parton to the particle level in terms of their absolute and relative size. Only the $t\bar{t}$ modelling uncertainties (except PS and hadronisation) are considered.

Name of the Uncertainty	Absolute Size	Relative Size [%]
FSR	0.0056	2.43
Recoil-to-top	0.0033	1.46
Top-quark mass	0.0032	1.41
Underlying event	0.0024	1.06
ME/PS matching (p_T^{hard})	0.0022	0.96
Colour Reconnection	0.0012	0.53
ISR	0.0007	0.33
Lineshape	0.0007	0.29
ISR h_{damp}	0.0005	0.22
pQCD $t\bar{t}$ NNLO QCD	0.0003	0.12

values are:

$$D_{\text{observed}} = -0.361 \pm 0.008[\text{stat.}]^{+0.034}_{-0.040}[\text{syst.}],$$
$$D_{\text{expected}} = -0.354 \pm 0.008[\text{stat.}]^{+0.033}_{-0.041}[\text{syst.}],$$

in the signal region of $m_{t\bar{t}} < 390\,\text{GeV}$,

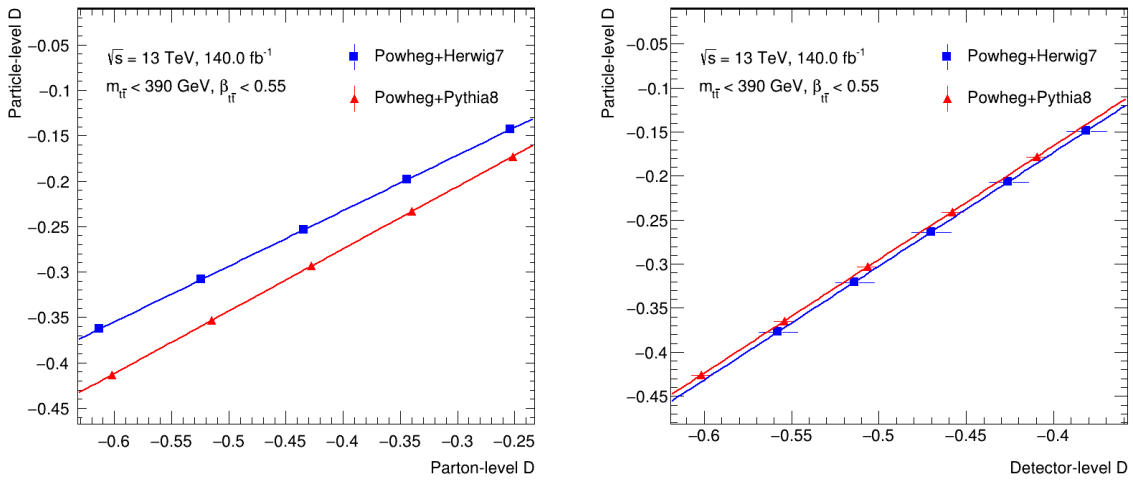
$$D_{\text{observed}} = -0.202 \pm 0.004[\text{stat.}]^{+0.031}_{-0.028}[\text{syst.}],$$
$$D_{\text{expected}} = -0.218 \pm 0.004[\text{stat.}]^{+0.031}_{-0.028}[\text{syst.}],$$

in the validation region of $390 < m_{t\bar{t}} < 500\,\text{GeV}$, and

$$D_{\text{observed}} = -0.128 \pm 0.003[\text{stat.}]^{+0.019}_{-0.018}[\text{syst.}],$$
$$D_{\text{expected}} = -0.146 \pm 0.003[\text{stat.}] \pm 0.019[\text{syst.}],$$

in the validation region of $m_{t\bar{t}} > 500\,\text{GeV}$.

The impacts of the systematic uncertainties on the observed value of D in the signal region remain similar to those in the case with quasi-bound-state effects. While the values in the validation regions remain largely unchanged, the values in the signal region differ slightly, with these differences discussed in detail in Section 7.11.3. The entanglement limits evaluated for both Powheg+Pythia and Powheg+Herwig are found to be identical to those obtained in the case with quasi-bound-state effects included.



(a) Particle-parton calibration curves.

(b) Detector-particle calibration curves.

Figure 7.13: Calibration curves mapping between different levels of event description for two generator setups: nominal POWHEG+PYTHIA (red) and alternative POWHEG+HERWIG (blue), used to estimate the PS and hadronisation uncertainty. A larger difference between the calibration curves is observed when mapping parton-level D to particle-level D.

7.11.3 Sensitivity to entanglement

The observed and expected values of the entanglement observable D, obtained with and without the inclusion of quasi-bound-state effects, $t\bar{t}_{\text{GFRW}}$, are summarised in Figure 7.14. The entanglement limits, propagated from parton to particle level for both generators, are also displayed. In the signal region, both the observed and expected values lie well below the entanglement limits, whereas this is not the case in the validation regions. This indicates that the $t\bar{t}$ system is ***maximally entangled near the $m_{t\bar{t}}$ threshold***, with the effect becoming diluted at higher $m_{t\bar{t}}$. This behaviour is consistent with theoretical expectations (Section 3.2.3), where $t\bar{t}$ pairs are predominantly produced in a spin-singlet configuration at low $m_{t\bar{t}}$, with a reduced contribution of this state beyond this region. Sensitivity to entanglement is expected to re-emerge at high $m_{t\bar{t}}$, where spin-triplet configurations dominate and can be probed with dedicated observables (Chapter 8).

The HERWIG predictions are shown without associated uncertainties and are included only to illustrate the level of agreement with PYTHIA. In the validation regions, where both predictions lie above the entanglement threshold, good agreement is observed. In contrast, a noticeable tension appears in the signal region, highlighting limitations in the current PS and hadronisation modelling at low $m_{t\bar{t}}$.

The observed values, both with and without $t\bar{t}_{\text{GFRW}}$, are consistent with the expected values within uncertainties. As expected, no significant difference is seen between the two scenarios in the validation regions, where quasi-bound-state effects are negligible. In the signal region, however, a small deviation is observed. In the model without $t\bar{t}_{\text{GFRW}}$ included, the data shows a higher degree of entanglement than the corresponding prediction. On the other hand, for the model with $t\bar{t}_{\text{GFRW}}$ included, the data shows a lower degree of entanglement than the corresponding prediction. This may suggest a smaller

contribution from quasi-bound-state effects in the data than predicted in the nominal model, although the deviation is very small and could also be due to statistical fluctuations in the fit. Furthermore, quasi-bound-state effects are expected to be more pronounced in narrower $m_{t\bar{t}}$ intervals, such as those used in dedicated threshold analyses [21, 22]. In the present analysis, which has a signal region that encompasses a wider $m_{t\bar{t}}$ range, their overall impact on the results is therefore reduced. The uncertainties on both the expected and observed values are comparable and larger than the agreement between data and prediction. They are primarily driven by uncertainties associated with the pQCD modelling of $t\bar{t}$ production.

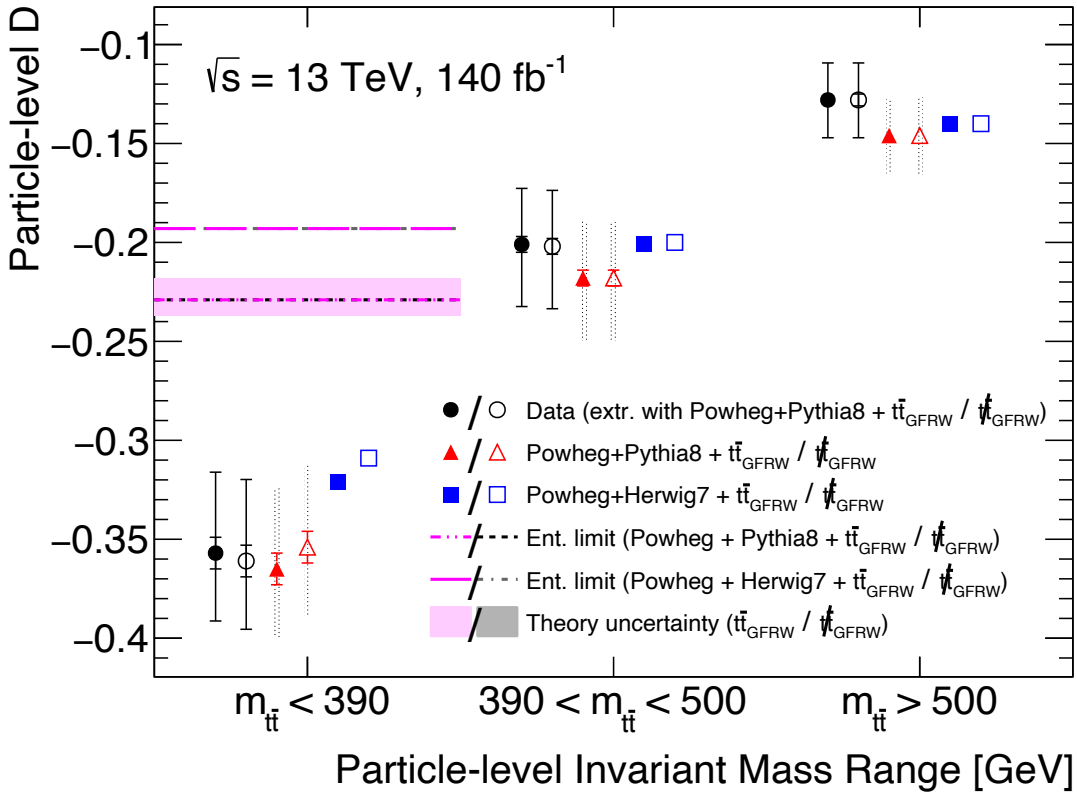


Figure 7.14: Overview of the values of the entanglement observable D at the particle level in the signal and the two validation regions. Filled markers correspond to the nominal fit model, which include quasi-bound-state effects, while open markers represent the scenario without quasi-bound-state effects. The observed values, extracted using the calibration curve of Powheg+Pythia with or without $t\bar{t}_{\mathrm{GFRW}}$ (without denoted by $\not{t\bar{t}}_{\mathrm{GFRW}}$), are shown as black circles. The expected values from Powheg+Pythia and Powheg+Herwig are indicated by red triangles and blue squares, respectively. For the data points, the inner error bars denote the statistical uncertainties, while the outer error bars represent the total uncertainties. For the Powheg+Pythia prediction (red triangles), the corresponding uncertainties are indicated by inner error bars and a shaded band, respectively. The entanglement limits, propagated from the parton to the particle level for both generators, are also shown together with their associated uncertainties.

Using the results described above, the entanglement sensitivity in the signal region is quantified by evaluating the significance, defined as the number of standard deviations by which the measured value

deviates from the entanglement limit. In this analysis, it is expressed as a ***Z-score***. For an observed or expected value at particle level, D_{particle}, the Z-score is given by:

$$Z = \frac{D_{\text{particle}} - D_{\text{particle}}^{-1/3}}{\sigma_D}, \tag{7.6}$$

Here, $D_{\text{particle}}^{-1/3}$ is the particle-level value corresponding to the parton-level entanglement limit of $-1/3$, obtained via the parton-to-particle level calibration curve. The quantity σ_D denotes the total uncertainty on D_{particle}, including both statistical and systematic contributions. The uncertainty is taken in the direction that maximises the deviation from the entanglement limit, $D_{\text{particle}}^{-1/3}$. The Z-test assumes that the distribution of D is approximately Gaussian and that the associated uncertainties are symmetric.

The central observed (expected) value includes the systematic uncertainty associated with PS and hadronisation modelling, while the entanglement limit does not. Since no dedicated systematic variations are available for Herwig, the relative impact of the systematic variations evaluated using Pythia is applied to the Herwig prediction. The observed (expected) significance is therefore evaluated twice: once with respect to the entanglement limit derived using Powheg+Pythia, and once using the limit obtained with Powheg+Herwig, while in both cases using the same central observed (expected) value. The observed (expected) significance in the signal region with respect to both parton shower generators are found to be:

$$Z(\text{Pythia} + t\bar{t}_{\text{GFRW}}) = 3.73\,(3.97)\,\sigma,$$
$$Z(\text{Herwig} + t\bar{t}_{\text{GFRW}}) = 4.78\,(5.01)\,\sigma.$$

The corresponding significances for the case where the signal is pQCD $t\bar{t}$ only is evaluated to be:

$$Z(\text{Pythia}) = 3.83\,(3.63)\,\sigma,$$
$$Z(\text{Herwig}) = 4.87\,(4.68)\,\sigma.$$

7.11.4 Results using the simplified model of quasi-bound-state effects

The calibration curve approach is also tested using the simplified model, η_t, instead of $t\bar{t}_{\text{GFRW}}$. In this calibration, the modelling uncertainties associated with η_t, as discussed in Section 7.9.3, are also taken into account. The particle-level values thus obtained are:

$$D_{\text{observed}} = -0.358 \pm 0.008[\text{stat.}]^{+0.033}_{-0.040}[\text{syst.}],$$
$$D_{\text{expected}} = -0.365 \pm 0.008[\text{stat.}]^{+0.033}_{-0.040}[\text{syst.}],$$

in the signal region of $m_{t\bar{t}} < 390\,\text{GeV}$,

$$D_{\text{observed}} = -0.201 \pm 0.004[\text{stat.}]^{+0.031}_{-0.028}[\text{syst.}],$$
$$D_{\text{expected}} = -0.219 \pm 0.004[\text{stat.}]^{+0.031}_{-0.028}[\text{syst.}],$$

in the validation region of $390 < m_{t\bar{t}} < 500\,\text{GeV}$, and

$$D_{\text{observed}} = -0.128 \pm 0.003[\text{stat.}]^{+0.019}_{-0.018}[\text{syst.}],$$
$$D_{\text{expected}} = -0.146 \pm 0.003[\text{stat.}] \pm 0.019[\text{syst.}],$$

in the validation region of $m_{t\bar{t}} > 500\,\text{GeV}$.

The entanglement limits evaluated for both Powheg+Pythia and Powheg+Herwig are calculated to be $-0.230^{+0.008}_{-0.011}[\text{syst.}]$ and -0.194, respectively.

Comparing these results with those in Section 7.11.1, it can be seen that the central values of D remain unchanged, or are very similar, to those obtained when quasi-bound-state effects are modelled using the $t\bar{t}_{\text{GFRW}}$ model. This indicates that, after object reconstruction, event selection, and region definitions, the signal region has limited sensitivity to differences between the two model implementations of quasi-bound-state effects. Furthermore, the overall systematic uncertainty remains unchanged, indicating a negligible impact from the modelling uncertainties of η_t. Table 7.12 summarises the impacts of the most dominant uncertainties on the observed value in the signal region. The impacts of the modelling uncertainties of η_t are also shown separately for comparison, illustrating that their contributions are significantly smaller than those of the leading systematic uncertainties. Since both the η_t and $t\bar{t}_{\text{GFRW}}$ models yield consistent central values, they are expected to exhibit similar sensitivities to modelling variations, leading to comparable systematic uncertainties. Such differences between central values and associated uncertainties between the two models are expected to become more pronounced when moving to a finer $m_{t\bar{t}}$ binning, where quasi-bound-state effects are more strongly enhanced.

Table 7.12: Ranking of the largest sources of systematic uncertainty in terms of their absolute and relative impact on the observed value of D after calibration from detector to particle level, using the signal model that employs the simplified model, η_t, for quasi-bound-state effects. Impact of the η_t modelling uncertainties are also listed for comparison.

Name of the Uncertainty	Absolute Size	Relative Size [%]
ME/PS matching (p_T^{hard})	0.0177	4.95
FSR	0.0106	2.97
Jet pile-up ρ topology	0.0097	2.72
Jet JER NP-1	0.0094	2.61
Parton Shower	0.0076	2.13
Jet JER NP-2	0.0061	1.70
Jet JER Data vs. MC	0.0060	1.69
Colour reconnection	0.0058	1.62
b-tag (Light) NP-0	0.0053	1.49
Jet JER NP-6	0.0051	1.42
η_t FSR	0.0004	0.12
η_t ISR	0.0003	0.08
η_t Parton Shower	0.0002	0.05

7.11.5 Comparison with other LHC entanglement measurements

As briefly discussed in Section 3.2.3, the first measurement of entanglement at the LHC was performed by the ATLAS experiment in the dilepton final state [17]. The corresponding results are shown in Figure 7.15. Entanglement was observed with a significance exceeding 5σ, which is substantially higher than that obtained in the present lepton+jets analysis. Although the reconstruction of the dilepton $t\bar{t}$ system is complicated due to the presence of two undetected neutrinos, this channel benefits from clean and well-identified spin analysers in the form of charged leptons. In contrast, the lepton+jets channel relies on the identification of the down-type jet as the hadronic spin analyser, which is experimentally more challenging. Additionally, the larger jet multiplicity in the lepton+jets final state further complicates the event reconstruction and introduces additional jet-related systematic uncertainties. Therefore, despite its smaller branching ratio, the dilepton channel provides superior sensitivity to entanglement. Nevertheless, the measurement presented in this thesis offers a novel aspect in that it provides access to complementary sources of systematic uncertainties not probed in the dilepton channel. It is also the first ATLAS entanglement measurement to include quasi-bound-state effects.

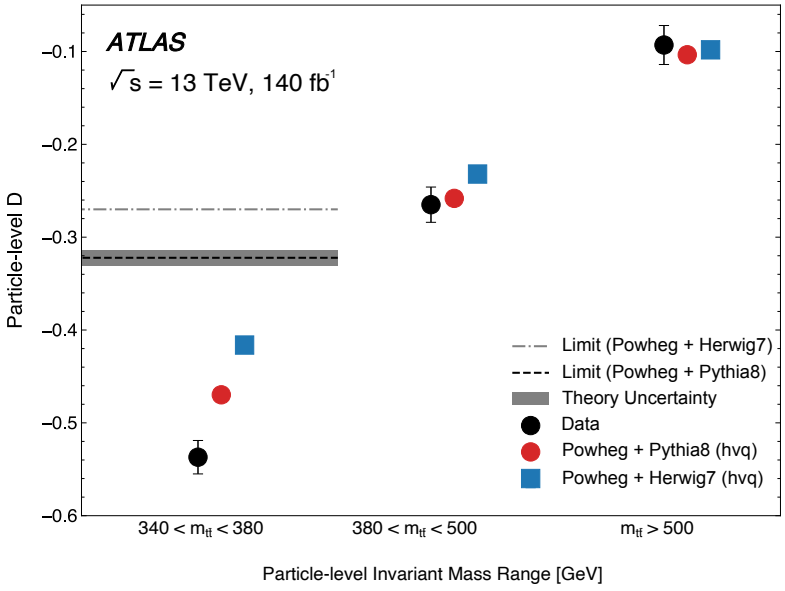


Figure 7.15: Overview of the ATLAS dilepton entanglement analysis result. The particle-level D results in the signal and validation regions are compared for various MC models. The entanglement limit, converted from the parton-level value $D = -1/3$ to the particle level, is shown [17].

The CMS result in the dilepton channel is shown in Figure 7.16. In the measurement, CMS considers a slightly different version[4], of the simplified model η_t [23]. It should also be noted that the definition of the low-$m_{t\bar{t}}$ region differs from that used in the ATLAS dilepton measurement, and additionally includes a selection on the longitudinal component of the $t\bar{t}$ boost, $\beta_z(t\bar{t})$. Furthermore, alongside the Powheg+Pythia prediction, an alternative sample is used to assess variations in the ME and

[4] The main difference with respect to the model used in this thesis is the resonance width: CMS uses a width of 7 GeV, compared to the 2.8 GeV considered here.

matching procedure. In this case, the $t\bar{t}$ sample is generated using MadGraph5_aMC@NLO, with ME jets matched to the parton shower using the FxFx prescription [254]. The value of D is extracted at the parton level via a binned profile likelihood fit to the $\cos\varphi$ distribution. While a significance exceeding 5σ is observed in this case as well, the observed degree of entanglement is smaller in the case where quasi-bound-state effects are included via the η_t model compared to the scenario in which it is neglected. This behaviour could also arise from just statistical fluctuations in the fits. Overall, the result is qualitatively consistent with observations in the present lepton+jets analysis, although a direct comparison is not straightforward due to differences in final states and phase-space definitions.

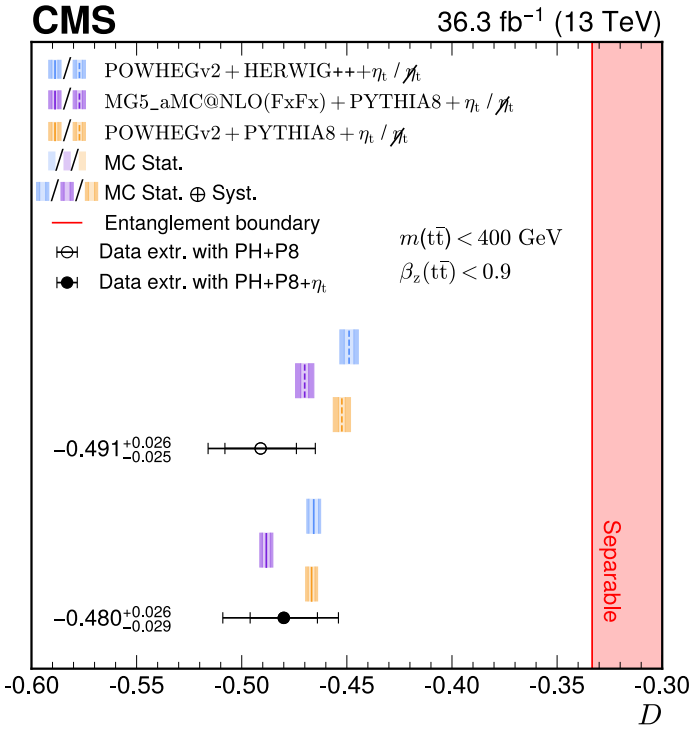


Figure 7.16: Summary of the CMS dilepton entanglement measurement. Shown are the observed parton-level D values (black points) compared with MC predictions with (solid) and without (dashed) the η_t state. Inner (outer) error bars show the statistical (total) uncertainty in data. Light (dark) shaded regions indicate the MC statistical (total) uncertainty. The entanglement limit $D = -1/3$ is shown as a shaded band [23]

CMS has also performed an entanglement measurement in the lepton+jets channel [24], providing the closest comparison to the present ATLAS lepton+jets result. The corresponding results are shown in Figure 7.17. The measurement of D is performed in the low-$m_{t\bar{t}}$ region, defined using a slight different $m_{t\bar{t}}$ selection of $m_{t\bar{t}} < 400$ GeV. Additionally, CMS performed the measurement in a phase-space that puts a requirement on the transverse momentum of the top quark, $p_T(t) < 50$ GeV. The right panel, corresponding to the low-$m_{t\bar{t}}$ region, is therefore most relevant for comparison with the analysis presented in this thesis. The quasi-bound-state effects are included via the simplified

η_t model, similar to the one used in this thesis[5]. However, the impact of these effects is treated as a systematic uncertainty rather than being incorporated directly into the $t\bar{t}$ signal model. For comparison with the data, predictions are provided both with and without these effects, along with several alternative modelling variations. The observed degree of entanglement is found to be lower than the predicted values, resulting in a significance of 1.6σ. In comparison, the ATLAS lepton+jets analysis presented in this thesis exhibits a higher sensitivity to entanglement effects. This is likely related to the tighter selection on $m_{t\bar{t}}$ applied in this analysis and the additional requirement on the boost of the $t\bar{t}$ system, $\beta_{t\bar{t}}$, which increases the fraction of gg-initiated events. Since gg production is expected to yield a larger fraction of entangled states at threshold compared to $q\bar{q}$ production, the selection on $\beta_{t\bar{t}}$ improves the overall sensitivity to $t\bar{t}$ entanglement.

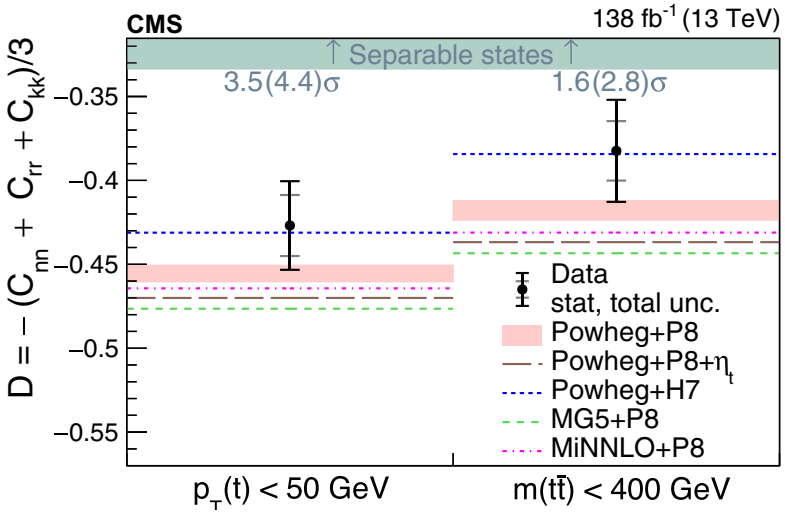


Figure 7.17: Measurement results of the entanglement observable D by CMS in the lepton+jets threshold region. The observed values are shown by black filled circles with the inner (outer) error bars showing statistical (total) uncertainty. The observed values are compared to Powheg+Pythia, Powheg+Pythia+η_t, Powheg+Herwig, MadGraph5_aMC@NLO+Pythia, and MiNNLO+Pythia predictions. Observed (expected) significance relative to separable states (green band) is given in units of standard deviations σ. The boundary of the green band corresponds to the entanglement threshold $D = -1/3$, below which the $t\bar{t}$ system is entangled [24].

[5] Width is 2.8 GeV.

CHAPTER 8

Probing entanglement with boosted top quarks

"The hole gets smaller, but you never fill it."

— Arcane 2

Following the observation of entanglement in the threshold region, this chapter turns to the high $m_{t\bar{t}}$ regime, where entanglement is also expected to be prominent in $t\bar{t}$ pairs. This region is of particular interest, as the top quarks are produced with large Lorentz boosts and can therefore become space-like separated before decaying. Therefore, it provides a unique and complementary environment to study quantum correlations in the top-quark sector. The investigation of quantum observables in this region also sets the stage for future measurements of Bell-like correlated states, which are expected to violate Bell inequalities if a loophole-free Bell test was experimentally achievable at colliders. Beyond their fundamental significance, these studies also offer a novel precision probe of the SM, enabling its predictions to be tested against observed data in an extreme kinematic region.

As discussed previously in Chapter 6, the measurement in the boosted region is a more recent analysis effort than its threshold counterpart and therefore benefits from improved reconstruction techniques, updated calibrations, and modern flavour-tagging techniques available in the ATLAS software framework for Run-3 data analyses. In particular, the identification of the down-type jet originating from the W boson of the hadronically decaying top quark relies on the state-of-the-art GN2 flavour-tagging algorithm[1]. At the same time, the boosted topology introduces additional experimental challenges. In this kinematic regime, the decay products of the hadronically decaying top quark tend to become increasingly collimated, such that the resulting jets may overlap and become difficult to resolve individually. As a result, the correct assignment of reconstructed physics objects to the underlying partons becomes significantly more challenging than in the threshold region. Furthermore, moving to very high $m_{t\bar{t}}$ phase space implies a rapidly falling production cross-section, and the analysis therefore becomes increasingly statistically limited. For this reason, in addition to the down-type jet, alternative spin analysers for the hadronically decaying top quark are investigated. The two analysers considered in this work are the ***optimal hadronic polarimeter*** and the ***W boson***. Their definitions are given in

[1] As this analysis is still under active development and is among the first to employ the new analysis tools, in particular the TopCPToolkit framework, the dedicated down-type jet definition using direct c-tagging via GN2 flavour-tagging algorithm had not yet been fully implemented within the framework at the time of writing of this thesis.

Section 3.2.1. For the semi-leptonically decaying top quark, the ***charged lepton*** is again used as the spin analyser, owing to its maximal spin analysing power and clean experimental signature.

Although there are important differences between the threshold and boosted analyses, many aspects of the general methodology remain unchanged. To avoid unnecessary repetition, common ingredients such as object and event selections, reconstruction strategies, and systematic uncertainty definitions are not repeated in full detail in this chapter, but are instead referenced back to the corresponding sections of Chapter 7. Only those components that differ specifically for this analysis are discussed here in detail.

8.1 Analysis strategy

The goal of this measurement is to extract the entanglement-sensitive observable, D_3, using the mean of the $\cos\theta'$ distribution (Equation 3.11). The analysis begins with the selection of the required baseline objects (Section 8.2), followed by a dedicated event selection (Section 8.3) designed to enhance the fraction of $t\bar{t}$ events in the lepton+jets channel, with particular emphasis on kinematics expected in the boosted regime. The $t\bar{t}$ system is then reconstructed, and the relevant angle is constructed (Section 8.4). These reconstructed kinematic quantities are then subjected to additional selections that define the signal regions and improve the sensitivity to entanglement effects, as described in Section 8.5. The $\cos\theta'$ distributions in these regions are used to determine the value of D_3, which is then corrected to the parton level using the unfolding procedure described in Section 8.6. The various sources of systematic uncertainty incorporated in the measurement are described in Section 8.7. The kinematic distributions entering the fitting procedure are discussed in Section 8.8. The results of the measurement and their interpretation are presented in Section 8.9.

Since the analysis strategy was still under development at the time of writing this thesis, with several aspects such as validation regions, systematic uncertainties, and non-prompt/fake background estimation still being finalised, the results presented in this thesis are *blinded*. This means that the entanglement observable is not extracted using collision data. Instead, the results are quoted using Asimov data, described in Section 8.8. This standard practice is employed to avoid potential bias from the observed data during the optimisation and validation of the measurement strategy.

8.2 Baseline object selection

8.2.1 Leptons

Selected electrons are required to satisfy the *Tight LH* identification WP. To ensure isolation, the *Tight_VarRad* isolation WP is applied, which combines track- and calorimeter-based isolation requirements (Section 6.2). In particular, the track-based component uses a variable-size cone around the electron, the radius of which decreases with increasing electron transverse momentum, while the calorimeter-based isolation is evaluated using a fixed-radius cone. The decreasing cone size reflects the fact that the decay products of highly boosted particles become increasingly collimated. This feature is particularly beneficial in the boosted topology considered in this analysis, where electrons are more likely to overlap with nearby hadronic activity originating from the boosted top-quark decay products. In such cases, a fixed-size cone would lead to a loss of efficiency due to the inclusion of

unrelated nearby tracks, whereas the variable-radius approach maintains higher signal efficiency while preserving background rejection.

Electrons are furthermore required to have a transverse momentum $p_T > 27$ GeV and the energy cluster to fall within the pseudorapidity range $|\eta_{\text{cluster}}| < 2.5$, excluding the transition region between the barrel and endcap calorimeters ($1.37 < |\eta_{\text{cluster}}| < 1.52$).

As for muons, the *Medium* identification quality criteria are applied, while isolation is based on the *Tight_VarRad* WP. The selected candidates are required to have a transverse momentum $p_{\text{T}} > 27$ GeV and a pseudorapidity $|\eta| < 2.5$. The requirements for the electron and muon selections are summarised in Table 8.1.

Table 8.1: Selection criteria for electrons and muons.

Feature	**Electron**	**Muon**
Pseudorapidity range	$\|\eta_{\text{cluster}}\| < 1.37$ or $1.52 < \|\eta_{\text{cluster}}\| < 2.5$	$\|\eta\| < 2.5$
Transverse momentum	$p_T > 27$ GeV	
Identification	Tight LH	Medium
Isolation	Tight_VarRad	

8.2.2 Jets

The boosted kinematic regime introduces challenges in the experimental reconstruction of jets originating from the decay products of the top quarks. While the semi-leptonically decaying top quark is comparatively straightforward to reconstruct, since its hadronic decay products consist of a single b-quark, the hadronically decaying top quark produces three partons in the final state which are detected experimentally as jets. In the boosted regime, these jets become highly collimated and may not be reconstructed as distinct jets.

Therefore, in this regime small-R jets are used to reconstruct the b-jet originating from the semi-leptonically decaying top quark. These jets are reconstructed using the anti-k_T jet clustering algorithm with a radius parameter of $R = 0.4$, using PFlow objects as inputs. Jet candidates are required to satisfy $p_{\text{T}} > 25$ GeV and $|\eta| < 2.5$. To further mitigate pile-up effects, a JVT requirement is applied to jets with $p_T < 60$ GeV and $|\eta| < 2.4$.

On the other hand, to capture the collimated decay products of the hadronically decaying top quark, reclustering is performed to construct RC jets as outlined in Section 6.4.4. The resulting RC jets are required to have $p_{\text{T}} > 300$ GeV and $|\eta| < 2.5$ in order to select well-reconstructed boosted top-quark candidates. In addition, they are required to contain two or three constituent jets in order to suppress the impact of JER effects, which are particularly significant for jets composed of a single constituent. The full set of selection requirements applied to jets in this analysis is summarised in Table 8.2.

8.3 Event selection

Following the baseline object selection, additional requirements are imposed to select $t\bar{t}$ events in the lepton+jets decay channel in the boosted regime. An initial set of preselection criteria, analogous to those used in the threshold entanglement analysis (Section 7.3.1), is applied. Furthermore, events are

Table 8.2: Selection criteria for small-R and RC jets.

Feature	Criteria
Small-R jets	
Pseudorapidity range	$\|\eta\| < 2.5$
Transverse momentum	$p_T > 25$ GeV
Pile-up suppression	JVT> 0.20 for $p_T < 60$ GeV and $\|\eta\| < 2.4$
RC jets	
Pseudorapidity range	$\|\eta\| < 2.0$
Transverse momentum	$p_T > 300$ GeV
Number of constituents	2 or 3

required to pass a specific set of single-lepton triggers, identical to those described in Section 7.3.2, except for the 2015 data-taking period, where a lower muon p_T threshold of 40 GeV is used for the non-isolated trigger instead of 50 GeV. This modification extends the efficiency of the non-isolated trigger into a low-p_T region, thereby increasing the acceptance of events containing muons that could fail the isolation requirement. This is particularly relevant in the boosted regime, where leptons are more likely to overlap with hadronic activity and fail standard isolation criteria. As a result, the choice improves the overall signal efficiency at the cost of a slightly higher reliance on non-isolated trigger selections at lower transverse momenta.

After the trigger selection, further object-level requirements are applied to select events consistent with $t\bar{t}$ production in the lepton+jets channel. Events are required to contain exactly one electron or muon satisfying the criteria defined in Section 8.2, with $p_{\mathrm{T}} > 27$ GeV. This requirement is chosen to select events from the region where the trigger efficiency is approximately constant, thereby ensuring full trigger efficiency. In addition, events must contain at least three small-R jets fulfilling the requirements detailed in Section 8.2. Of these, at least two must be b-tagged using the GN2 algorithm at the 85% efficiency WP (Section 6.4.6). The conventional requirement of at least four small-R jets is relaxed in this analysis, as the boosted topology increases the likelihood of jet merging due to the collimation of decay products. This selection therefore retains events in which two jets are reconstructed as a single jet. Finally, to identify the hadronically decaying top quark, events are required to contain at least one RC jet satisfying the criteria described in Section 8.2, with mass $m_{\mathrm{jet}}^{\mathrm{RC}}$ in the range $120 < m_{\mathrm{jet}}^{\mathrm{RC}} < 200$ GeV. This mass requirement is made to ensure that the RC jet is most likely originating from the hadronically decaying top quark. A summary of all the event selection criteria is provided in Table 8.3.

8.4 Event reconstruction

For events passing the above selection criteria, the full $t\bar{t}$ system is reconstructed. The reconstruction is performed analytically within the TopCPToolkit framework. The procedure is as follows. If an event contains exactly one RC jet, the jet is identified as the hadronically decaying top quark candidate. If multiple RC jets are present, the jet with mass closest to the top-quark mass, $m_t = 172.5$ GeV, is selected. The chosen RC jet is further required to satisfy a distance criterion with respect to the selected charged lepton (l), $\Delta R(\text{RC jet}, l) > 1.0$ and to contain exactly one b-tagged constituent jet.

Table 8.3: Summary of the event selection criteria (Boosted analysis).

Selection	Criteria
Data quality criteria	Data events in GRLs (2015-2018); jet cleaning, primary vertex selection, removal of LAr noise bursts, tile and SCT errors.
Single lepton trigger	Separate single-electron or single-muon triggers
Exactly one lepton	== 1 e or μ with $p_T > 27$ GeV
Small-R jets	≥ 3 jets, $p_T > 25$ GeV
RC jets	≥ 1 jet(s), $p_T > 300$ GeV, 2 or 3 constituents, $120 < m_{\text{jet}}^{\text{RC}} < 200$ GeV
b-tagging	≥ 2 b-tagged jets at 85% WP of GN2

This distance requirement ensures a clear separation between the hadronically decaying top-quark candidate and the lepton originating from the semi-leptonically decaying top quark. Among the selected small-R jets, the jet closest to the lepton such that $\Delta R(\text{jet}, l) \leq 2.0$ is assigned as the b-jet originating from the semi-leptonically decaying top quark. In addition, the invariant mass of the lepton and the b-jet is required to satisfy $m(l, b) < 180$ GeV. This invariant mass requirement exploits the kinematic upper limit of $t \to bl\nu$ decay to suppress incorrect lepton-jet pairings, which tend to produce larger $m(l, b)$ values. The selected objects are then used to analytically reconstruct the four-momentum of the neutrino originating from the semi-leptonically decaying top quark, as described earlier in Section 7.4.2. The four-momenta of the hadronic and semi-leptonic top quarks are subsequently reconstructed from their respective decay products, and the invariant mass of the $t\bar{t}$ system is computed from these.

Definition of the angular observable

The angle $\cos\theta'$, used for the measurement of entanglement in the boosted regime (Section 3.2.3), is constructed analogously to $\cos\varphi$, using the directions of the two spin analysers associated with the semi-leptonically and hadronically decaying top quarks, respectively, with a slight modification. The spin analysers are first boosted into the respective parent top-quark rest frames, following the two-step procedure described in Section 7.4.2. In these frames, their three-momenta are expressed in the helicity basis. As explained in Section 3.2.2, the helicity basis provides good sensitivity to the dominant spin configurations in the high $m_{t\bar{t}}$ region, where opposite-helicity $t\bar{t}$ states dominate.

The resulting three-momenta of the two spin analysers in the helicity basis are then normalised. To construct the angle $\cos\theta'$ sensitive to the observable D_3, a reflection transformation is applied to one of the spin analysers in the helicity basis, such that the signs of its r- and k-components are reversed while the n-component remains unchanged[2]. The angle $\cos\theta'$ is then defined as the scalar product of the resultant unit vectors of the two spin analysers.

[2] In some references, the transformation used to construct $\cos\theta'$ is implemented by flipping the sign of the n-component rather than the r- and k-components. This corresponds to a different convention for the definition of the helicity basis or the observable D_3. The underlying physics remains unchanged, but the overall sign convention differs.

8.5 Signal regions definition

The signal regions in this analysis are designed to be enriched in events where the two top quarks are highly Lorentz-boosted. As the Lorentz boost increases, the decay products of the hadronically decaying top quark become increasingly collimated. At sufficiently high boost, this collimation can cause two or more partons to become unresolved, resulting in a reduced number of reconstructed subjets or constituents within the RC jet associated with the hadronically decaying top quark.

Taking this kinematic feature into account, two orthogonal signal regions are defined at the reconstruction level, corresponding to RC jets with either two or three constituents. In both regions, the charged lepton is used as the spin analyser of the semi-leptonically decaying top quark. For the hadronically decaying top quark, the choice of spin analyser depends on the number of constituents in the RC jet. For events with three constituents, the decay topology is sufficiently resolved to construct the optimal hadronic polarimeter, using the two non-b-tagged constituent jets, following the definition introduced in Section 3.2.1. This region is denoted as ***SR1***. For events with two constituents, the topology typically corresponds to a more boosted configuration in which two of the hadronic decay jets are merged and cannot be resolved separately. In this case, one of the constituents may correspond to the hadronically decaying W boson. Therefore, in this region, the reconstructed W boson is used as the hadronic spin analyser. This region is denoted as ***SR2***. It is worth noting that, in this region, the merged constituents may not always originate from the two light-quark decay products of the W boson. In rare cases, the b-jet and one of the decay products of the W boson may be clustered into the same constituent. However, this contribution is expected to be small in the kinematic regime considered. To enhance sensitivity to quantum entanglement in the boosted regime, both signal regions are required to satisfy $m_{t\bar{t}} > 800$ GeV and $|\cos\Theta| < 0.4$, where Θ is the angle between the top quark direction and the beam axis.

Unlike the threshold entanglement measurement, the reconstruction-level value of D_3 is unfolded directly to the parton level using the procedure described in Section 8.6. The same kinematic requirements, $m_{t\bar{t}} > 800$ GeV and $|\cos\Theta| < 0.4$, are applied at parton level to define the boosted fiducial phase space. For both reconstruction-level signal regions, a common parton-level definition of the hadronic spin analyser is used, namely the optimal hadronic polarimeter. Therefore, two different reconstruction-level spin analyser definitions (the optimal polarimeter in SR1 and the W boson in SR2) are unfolded to a single common parton-level observable. This approach provides a consistent observable that is independent of the topology at detector level, allowing both signal regions to constrain the same distribution and to be combined naturally in a single unfolding fit. Since the decay products are uniquely defined at parton level and are not subject to jet merging or reconstruction ambiguities, the optimal hadronic polarimeter is the most natural hadronic spin analyser choice in this unfolding fit setup[3]. Accordingly, a spin analysing power of $\alpha = 0.67$[4] is used for the hadronic spin analyser, together with $\alpha = 1$ for the charged lepton. Each reconstruction-level signal region is further subdivided according to the flavour of the charged lepton, yielding separate electron (e+jets) and muon (μ+jets) regions. This categorisation increases the number of statistically independent regions

[3] In principle, the down-type quark would be the optimal choice of spin analyser. However, since this definition is not yet implemented at the detector or reconstruction level, an alternative analyser corresponding to one of the reconstruction-level signal regions is used.

[4] The value $\alpha = 0.67$ corresponds to the spin analysing power used in the analysis at the time of its development, rather than the currently adopted value of $\alpha = 0.64$. Therefore, the expected results presented in this thesis differ from those obtained with the updated value. However, the choice of α is not expected to affect the relative impact of the uncertainties.

entering the fit, with complementary systematic uncertainties, and provides additional flexibility to the fit parameters. In total, four reconstruction-level regions are defined: SR1 e+jets, SR1 μ+jets, SR2 e+jets, and SR2 μ+jets. A schematic overview of the region definitions and their mapping to the common parton-level observable is shown in Figure 8.1.

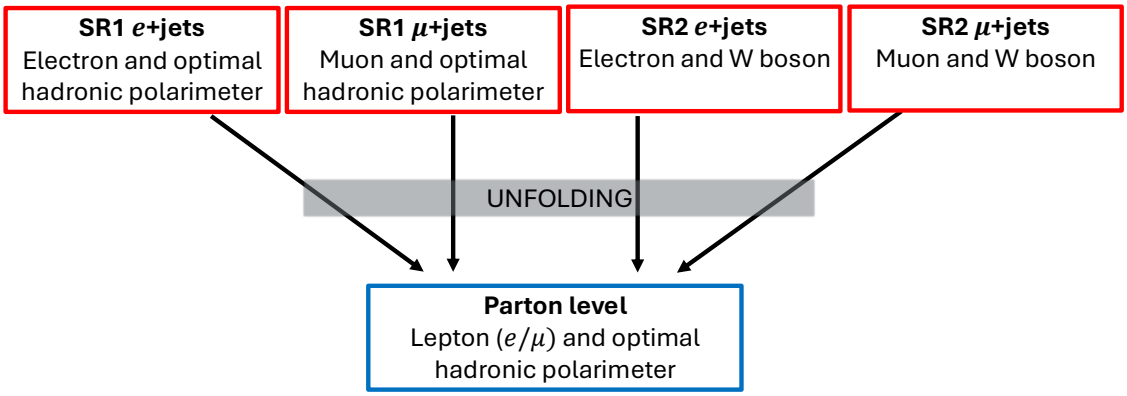


Figure 8.1: Illustration of the region definitions used for unfolding at the reconstruction level (top, shown in red) and at the parton level (bottom, shown in blue). All regions at both levels are defined with the selections $m_{t\bar{t}} > 800$ GeV and $|\cos\Theta| < 0.4$.

8.6 Profile-likelihood unfolding

As discussed in Section 3.2.3, the condition $D_3 < -1/3$ is strictly defined at the parton level. Therefore, an unfolding procedure is required to relate the observable measured at the detector level to its corresponding parton-level value. For this measurement, a ***profile-likelihood unfolding (PLU)*** method, implemented in the TRexFitter framework [239], is employed. In this approach, the unfolding is formulated in such a way that it can be treated as a standard profile-likelihood fit, which is a fitting procedure used to determine the set of parameters that best describe the observed data under a given expectation model. In other words, instead of performing unfolding as a separate step, it is incorporated directly into the fitting procedure used to extract the observable. This allows the parton-level value of D_3 to be determined in a consistent way within the fit itself, which motivates the term *profile-likelihood unfolding*.

In this section, the methodology underlying the PLU approach is described. It begins with a brief introduction to profile-likelihood fits, followed by an overview of the traditional unfolding procedure, and concludes with a description of how these concepts are combined in the PLU procedure used in this analysis. Additional practical details specific to the unfolding and extraction of D_3 are also discussed.

8.6.1 Profile-likelihood fits

In statistical analyses involving binned histograms, such as this measurement, where n_i denotes the number of observed events in the i^{th} bin, one seeks to construct a prediction ν_i that depends on a set

of theoretical parameters $\vec{\alpha} = (\alpha_1, ..., \alpha_m)$, such that n_i is well described by ν_i. The probability of observing n_i events given an expectation ν_i is described by the Poisson distribution:

$$\mathcal{P}(n_i, \nu_i(\vec{\alpha})) = \frac{e^{-\nu_i} {\nu_i}^{n_i}}{n_i!}. \tag{8.1}$$

Using Equation 8.1, a function known as the ***likelihood***, which incorporates the information from all bins of the histogram, can be constructed as [255]:

$$L(\vec{\alpha}) = \prod_{i=1}^{N_{\text{bins}}} P(n_i|\nu_i(\vec{\alpha})), \tag{8.2}$$

where N_{bins} denotes the total number of bins in the histogram. Sometimes, more than one histogram are used to define the different signal regions, as in this analysis. In such cases, the likelihood is constructed as the product over all bins and all regions, and is given by:

$$L(\vec{\alpha}) = \prod_{r \in \text{regions}} \prod_{i=1}^{N_{\text{bins}}} P(n_{ri}|\nu_{ri}(\vec{\alpha})),$$

where r denotes the region, and i denotes the bin number within each region.

Maximising the likelihood

The likelihood function can be used to estimate the most probable values of the parameters that best describe the data. This is done using the method of *maximum-likelihood estimation*, which consists of maximising the likelihood with respect to $\vec{\alpha}$. In practice, it is more convenient to maximise the logarithm of the likelihood, $\ln L$, which transforms the product into a sum. For a given histogram (using Equation 8.2), it can be written as:

$$\ln L(\vec{\alpha}) = \sum_{i=1}^{N_{\text{bins}}} \ln P(n_i \mid \nu_i(\vec{\alpha})) \propto \sum_{i=1}^{N_{\text{bins}}} \left[n_i \ln \nu_i(\vec{\alpha}) - \nu_i(\vec{\alpha}) \right].$$

The best-fit parameter values, $\hat{\vec{\alpha}}$, are estimated by maximising this likelihood, which is equivalent to solving:

$$\left. \frac{\partial \ln L(\vec{\alpha})}{\partial \alpha_k} \right|_{\vec{\alpha}=\hat{\vec{\alpha}}} = 0 \quad \forall k.$$

Since the likelihood depends on the parameters $\vec{\alpha}$, correlations between different parameters (α_i, α_j) must be taken into account for a consistent estimation of their uncertainties.

Under the Gaussian approximation, these uncertainties can be evaluated using the ***covariance matrix***, $\hat{V}$, defined as the inverse of the Hessian of the log-likelihood at its maximum:

$$\hat{V}_{ij} = \left[- \left. \frac{\partial^2 \ln L}{\partial \alpha_i \, \partial \alpha_j} \right|_{\hat{\alpha}_i, \hat{\alpha}_j} \right]^{-1}, \tag{8.3}$$

where the indices i, j run over the number of parameters in the vector $\vec{\alpha}$. The uncertainty on a given parameter α_k is then given by

$$\sigma_{\alpha_k} = \sqrt{\hat{V}_{kk}}. \tag{8.4}$$

A more precise evaluation is obtained via a method that consists of identifying the α_k where the value of $\ln L$ is reduced by 1/2, and can lead to asymmetric uncertainties:

$$\ln L(\alpha_k \pm \sigma_{\alpha_k}^{\text{up/down}}) = \ln L(\hat{\alpha_k}) - \frac{1}{2}.$$

Parameters of interest and nuisance parameters

In practical applications, not all parameters in the likelihood function play the same role. The full set of parameters, denoted by $\vec{\alpha}$, can be separated into ***parameters of interest (POIs)***, which correspond to the physical quantities being measured, and ***nuisance parameters (NPs)***, which encode systematic uncertainties affecting the measurement. The expected number of events in bin i of a given histogram can therefore be written as:

$$\nu_i(\vec{\mu}, \vec{\theta}) = S_i(\vec{\mu}, \vec{\theta}) + B_i(\vec{\theta}),$$

where S_i and B_i denote the expected signal and background yields in bin i, respectively. The vector $\vec{\mu} = (\mu_1, \ldots, \mu_n)$ represents the POIs, while $\vec{\theta} = (\theta_1, \ldots, \theta_m)$ denotes the NPs. Therefore, in this analysis, for a given region at the reconstruction level, the likelihood function is constructed as a product of Poisson probabilities over all bins:

$$L(\vec{\mu}, \vec{\theta}) = \prod_{i=1}^{N_{\text{bins}}} \mathcal{P}\left(n_i \middle| S_i(\vec{\mu}, \vec{\theta}) + B_i(\vec{\theta})\right). \tag{8.5}$$

Systematic uncertainties are incorporated through the NPs in the likelihood model, allowing for variations in both the shape and normalisation of signal and background predictions to reflect experimental and theoretical uncertainties. Each NP is defined with a nominal value θ^0 and a corresponding $\pm\sigma$ variation $\Delta\theta$. The effect of each uncertainty on the predicted distributions is implemented by interpolating between the nominal and systematically varied templates as the parameter θ deviates from its nominal value $\theta^0 \pm \Delta\theta$. To constrain large deviations from prior knowledge, penalty terms are introduced for each NP. Different choices for functions describing these penalty terms are possible, with the most common choice being Gaussian penalty terms [255]:

$$L(\vec{\mu}, \vec{\theta}) = \prod_{i=1}^{N_{\text{bins}}} \mathcal{P}\left(n_i \middle| S_i(\vec{\mu}, \vec{\theta}) + B_i(\vec{\theta})\right) \prod_{j=1}^{m} \mathcal{G}(\theta_j \mid \theta_j^0, \Delta\theta_j),$$

where $\mathcal{G}$ is the Gaussian penalty term given by

$$\mathcal{G}(\theta_j \mid \theta_j^0, \Delta\theta_j) = \frac{1}{\sqrt{2\pi}\Delta\theta_j} \exp\left(-\frac{(\theta_j - \theta_j^0)^2}{2(\Delta\theta_j)^2}\right).$$

Depending on the nature of the systematic uncertainty and the physics process, NPs may be treated as fully correlated across bins and processes, or correlated across bins but independent between

processes. The likelihood fit is performed by maximising $L(\vec{\mu}, \vec{\theta})$ with respect to both the POIs and the NPs. In this procedure, the NPs are *profiled*, meaning they are adjusted to their best-fit values for each choice of $\vec{\mu}$, yielding estimators $\hat{\vec{\mu}}$ and $\hat{\vec{\theta}}$. These correspond to the values that provide the best simultaneous description of the observed data and the prior constraints. The behaviour of the NPs after the fit provides important diagnostic information. Their fitted values are often expressed in terms of ***pulls***, defined as:

$$\text{pull}(\theta_j) = \frac{\hat{\theta}_j - \theta_j^0}{\Delta\theta_j}.$$

Under this transformation, the Gaussian constraint becomes a standard normal distribution with mean 0 and standard deviation 1. In practice, this corresponds to a reparametrisation in which the NPs are expressed in units of their prior uncertainties, such that $\theta_j^0 = 0$ and $\Delta\theta_j = 1$ by construction. These parameters can be restricted in a fit to the data, resulting in post-fit uncertainties that are smaller than their prior values. Such NPs are referred to as ***constrained***. This indicates that the data provides significant information about the corresponding systematic effect. It can also mean that the uncertainty corresponding to the constrained NP is poorly or crudely defined. Non-zero pulls indicate that the data prefers values of NPs shifted away from their nominal expectations. Both large pulls and strong constraints should be carefully examined to assess the modelling of the associated systematic uncertainties and their impact on the fit.

Moreover, the likelihood model incorporates additional NPs that account for the finite statistical precision of MC simulated samples. These are often referred to as γ-*parameters* and act as bin-by-bin normalisation factors that scale the MC prediction. Such parameters are particularly important in bins with limited MC statistics, where fluctuations in the simulated event yields can lead to significant uncertainties in the predicted templates. In contrast to systematic uncertainties, these γ-parameters are treated as uncorrelated between bins, with one independent parameter per bin. To account for these effects, the likelihood is extended by introducing an auxiliary measurement via Poisson constraint terms. The contribution to the likelihood can then be written as [255]:

$$\prod_{b \in N_{\text{bins}}} \mathcal{P}\big(n_b \mid \nu_b(\vec{\alpha}) + \gamma_b \nu_b^{\text{MC}}(\vec{\alpha})\big)\mathcal{P}\big(m_b \mid \gamma_b \tau_b\big),$$

where n_b denotes the number of observed events in bin b. The expected yield is split into two contributions: $\nu_b(\vec{\alpha})$, corresponding to components with negligible statistical uncertainty, and $\nu_b^{\text{MC}}(\vec{\alpha})$, corresponding to components affected by finite MC statistics. Both terms depend on the full set of parameters $\vec{\alpha}$. The NP γ_b is a parameter that scales the MC prediction in bin b, accounting for the fact that the true expectation may differ from the nominal MC estimate due to statistical fluctuations. In this formulation, the auxiliary term $\mathcal{P}(m_b \mid \gamma_b \tau_b)$ constrains γ_b through an effective MC event count m_b, with expectation value τ_b. The parameter γ_b is thus centred at unity, corresponding to the nominal MC prediction. The statistical uncertainty in the MC prediction is encoded through the relation $m_b = \left(\frac{\nu_b^{\text{MC}}}{\delta_b}\right)^2$, where δ_b denotes the statistical uncertainty on the MC yield in bin b. This construction ensures that the allowed variations of γ_b are consistent with the statistical precision of the MC sample, with larger uncertainties leading to weaker constraints on γ_b.

8.6.2 Unfolding

The goal of unfolding is to correct for detector effects and determine the underlying distribution at the parton level from measurements performed at the reconstruction level. The parton-level distribution thus obtained is called the *unfolded distribution*. In principle, unfolding can also be performed to the particle level, which can provide a more robust definition in some cases. However, in this analysis the parton level is used.

Let the signal distribution at the parton level be described by a set of bin yields S_j^{parton}, where j runs over the bins of the parton-level distribution. In this framework, the POIs $\vec{\mu} = (\mu_1, \ldots, \mu_n)$ are defined as normalisation factors that scale the signal yield in each parton-level bin, where $n = N_{\text{bins}}^{\text{parton}}$ denotes the number of parton-level bins. This number need not be the same as the number of bins at the reconstruction level. The parton-level distribution can therefore be parameterised as:

$$S_j^{\text{parton}}(\vec{\mu}, \vec{\theta}) = \mu_j \, S_j^{\text{parton,nom}}(\vec{\theta}),$$

where $\vec{\theta}$ denotes the NPs and $S_j^{\text{parton,nom}}$ denotes the nominal parton-level bin content, corresponding to the expected prediction before the application of POI and NP effects.

If a mapping can be constructed that relates the parton-level distribution to the distribution at the reconstruction level, accounting for detector response and acceptance effects, then the expected signal yield in a reconstruction-level bin i can be expressed as a function of the parton-level bins:

$$S_i(\vec{\mu}, \vec{\theta}) = \sum_{j=1}^{N_{\text{bins}}^{\text{parton}}} \mathcal{R}_{ij}(\vec{\theta}) S_j^{\text{parton}}(\vec{\mu}, \vec{\theta}),$$

where $\mathcal{R}_{ij}$ denotes the ij^{th} component of a matrix $\mathcal{R}$ that encodes this mapping.

Substituting this expression into the likelihood at reconstruction level (ignoring constraint terms for simplicity) gives:

$$L(\vec{\mu}, \vec{\theta}) = \prod_{i=1}^{N_{\text{bins}}} \mathcal{P}\left(n_i \,\middle|\, \sum_{j=1}^{N_{\text{bins}}^{\text{parton}}} \mathcal{R}_{ij}(\vec{\theta}) \, S_j^{\text{parton}}(\vec{\mu}, \vec{\theta}) + B_i(\vec{\theta}) \right).$$

From the above, it follows that the best-fit values of $\vec{\mu}$ are obtained by fitting the detector-level prediction, constructed from the parton-level distribution via the $\mathcal{R}$ matrix, to the observed data. These parameters act as bin-by-bin normalisation factors for the parton-level distribution relative to a nominal reference. Applying the fitted values $\hat{\vec{\mu}}$ to the nominal parton-level spectrum yields the unfolded distribution:

$$\hat{S}_j^{\text{parton}} = \hat{\mu}_j \, S_j^{\text{parton,nom}}.$$

This procedure, in which the parton-level distribution is inferred through a likelihood fit with profiled NPs and POIs, is referred to as *profile-likelihood unfolding*.

The $\mathcal{R}$ matrix, that takes care of the mapping above, is commonly referred to as the ***response matrix***. It encodes the relationship between the parton-level and detector-level distributions, incorporating the effects of detector resolution, bin-to-bin migrations, and event selection. It is constructed from three main components, described below.

Migration matrix The migration matrix M accounts for detector smearing effects, where events originating in a given parton-level bin are reconstructed in different detector-level bins. An element M_{ij} represents the probability that an event generated in parton-level bin j is reconstructed in detector-level bin i.

Efficiency The efficiency histogram f^{eff} quantifies the fraction of events generated at parton-level that also pass the reconstruction-level selection. The efficiency in the j^{th} parton-level bin is defined as:

$$f_j^{\text{eff}} = \frac{N_j^{\text{reco}} \cap N_j^{\text{parton}}}{N_j^{\text{parton}}},$$

where N denotes the number of events, and the superscripts reco and parton refer to reconstruction (detector) and parton levels, respectively. Thus, f_j^{eff} represents the probability that an event generated in parton-level bin j passes the detector-level selection.

Acceptance The acceptance histogram f^{acc} quantifies the fraction of reconstructed events at the detector level that are also compatible with the parton-level selection. The acceptance in the i^{th} detector-level bin is defined as:

$$f_i^{\text{acc}} = \frac{N_i^{\text{reco}} \cap N_i^{\text{parton}}}{N_i^{\text{reco}}}.$$

It therefore represents the probability that an event reconstructed in bin i also satisfies the parton-level selection criteria.

Using these three components, an element $\mathcal{R}_{ij}$ of the response matrix is then obtained as:

$$\mathcal{R}_{ij} = \frac{1}{f_i^{\text{acc}}} M_{ij} f_j^{\text{eff}}.$$

Therefore, the full PLU procedure proceeds as follows. First, each bin of the nominal signal parton-level histogram, $S_j^{\text{parton,nom}}$, is mapped to its corresponding reconstruction-level distribution using the response matrix. This step is referred to as ***folding***. Figure 8.2 shows an example of folding a single parton-level bin to the reconstruction level. Depending on the amount of bin migration, the content of a given parton-level bin can be distributed across several reconstruction-level bins. This folding procedure is repeated for every parton-level bin. Since each bin is associated with a normalisation factor μ_j, the corresponding folded distribution is scaled by the same factor. The sum of all folded signal contributions gives the expected signal distribution at the reconstruction level. This signal prediction is then added to the expected background contribution to obtain the total predicted detector-level distribution. A profile-likelihood fit of the total detector-level prediction to the observed data determines the best-fit values $\hat{\hat{\mu}}$. These fitted scale factors are then applied to the nominal parton-level spectrum, yielding the unfolded distribution that best describes the measured data. The systematic uncertainties, encoded through the NPs, are incorporated by allowing the response matrix, signal templates, and background predictions to vary within the fit. For each systematic source, modified response matrices are constructed and the folding procedure is repeated during the fit.

In the unfolding procedure, detector smearing and limited acceptance often lead to strong correlations between neighbouring bins of the parton-level distribution. This effect is particularly pronounced

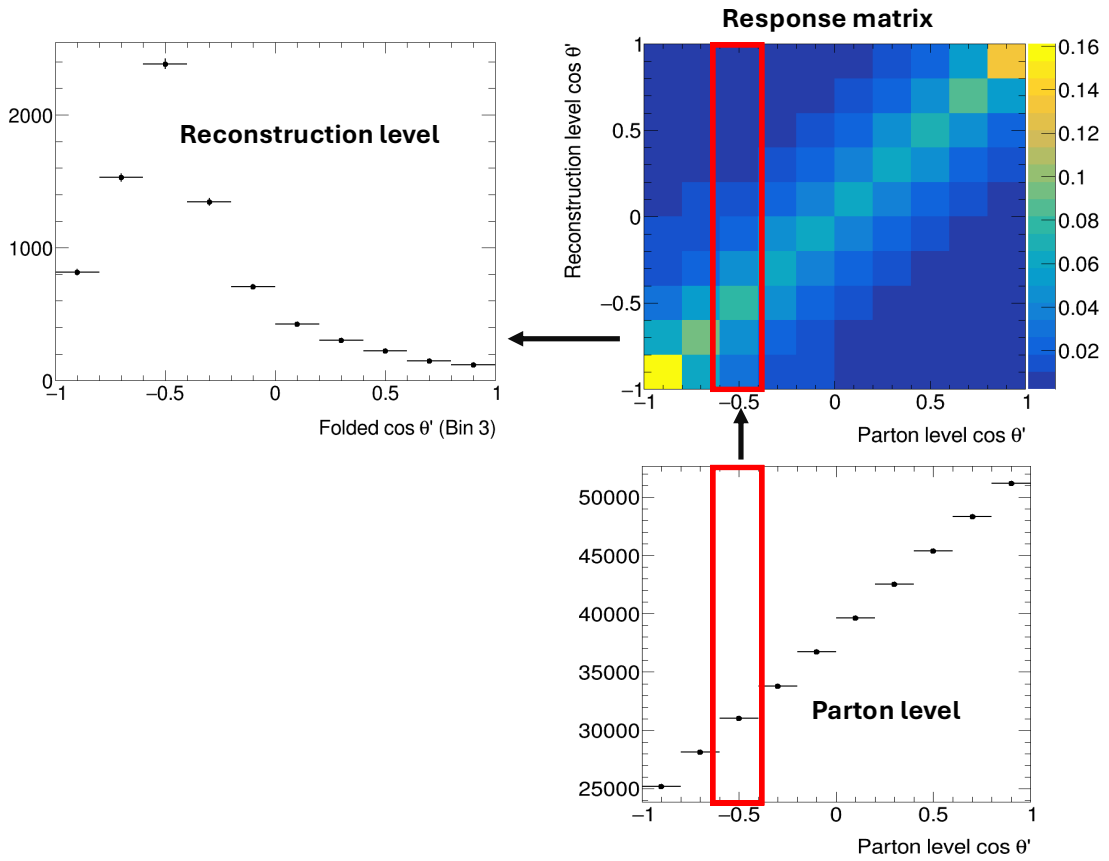


Figure 8.2: Schematic illustration of the folding procedure from the parton level to the reconstruction level using the response matrix. In this example, both the parton- and reconstruction-level distributions are divided into 10 bins, resulting in a 10×10 response matrix. The matrix is predominantly diagonal, indicating that most events are reconstructed in the same bin as defined at the parton level. For instance, the third parton-level bin (highlighted in red) maps primarily onto the third bin of the folded reconstruction-level distribution, with smaller contributions to neighbouring bins arising from bin-to-bin migrations.

when the migration matrix contains sizable off-diagonal elements, indicating significant bin-to-bin migrations between parton and reconstruction level. In such cases, statistical fluctuations in the measured data can be amplified, resulting in large oscillations or unstable bin-by-bin variations in the unfolded distribution. To suppress these effects and stabilise the fit, a procedure known as ***regularisation*** is introduced. It is implemented by adding an additional penalty term to the likelihood, thereby favouring solutions that are more stable or physically reasonable. Several regularisation schemes exist, one of the most widely used being *Tikhonov regularisation* [239].

In general, Tikhonov regularisation can be formulated in different ways depending on the choice of the constraint. The simplest form is the *zero-order* (reference-based) regularisation, $R_{\text{ref}}(\vec{\mu})$, in which each parameter is constrained to remain close to a reference value. For the normalisation factors μ_i, this corresponds to penalising deviations from the nominal prediction, which is taken to be unity in this analysis:

$$R_{\text{ref}}(\vec{\mu}) = \sum_{i=1}^{N_{\text{bins}}^{\text{parton}}} \left(\mu_i - 1\right)^2 .$$

This form of regularisation suppresses deviations from the nominal prediction on a bin-by-bin basis. While it provides strong stabilisation, it can also introduce a bias towards the reference shape, since the solution is explicitly pulled towards $\mu_i = 1$ in every bin.

A more flexible approach is to regularise the shape of the distribution rather than its absolute normalisation. In this analysis, Tikhonov regularisation is therefore used in its *second-derivative* (smoothness) form, $R_{\text{smooth}}(\vec{\mu})$, which penalises local curvature in the unfolded spectrum:

$$R_{\text{smooth}}(\vec{\mu}) = \sum_{i=2}^{N_{\text{bins}}^{\text{parton}}-1} \left(\mu_{i+1} - 2\mu_i + \mu_{i-1}\right)^2 .$$

This corresponds to a discretised second derivative and suppresses rapid bin-to-bin fluctuations, while not constraining the solution towards any specific reference shape.

The regularised likelihood is then written as:

$$L_{\text{reg}}(\vec{\mu}, \vec{\theta}) = L(\vec{\mu}, \vec{\theta}) \, \exp\left[-\frac{1}{2}\tau^2 \, R(\vec{\mu})\right],$$

where $L(\vec{\mu}, \vec{\theta})$ is the nominal profile likelihood and $R(\vec{\mu})$ denotes the chosen regularisation term. The parameter τ controls the strength of the constraint, determining the trade-off between statistical fluctuations and stability of the unfolded result. For $\tau \to 0$, the unfolding becomes fully data-driven and corresponds to the unregularised fit. Increasing τ suppresses fluctuations induced by limited statistics and strong bin migrations, but may reduce sensitivity to genuine fine structures in the distribution. In this analysis, the current fit configuration uses $\tau = 1$. Various other τ values have been tested, and $\tau = 1$ is currently used as it provides a good balance between the mitigation of bin-migration effects and the preservation of sensitivity to the underlying physics.

In principle, when regularisation is applied, the standard uncertainty estimates obtained from the Hessian matrix or from likelihood scans may no longer provide a complete description of the parameter uncertainties, since the regularisation introduces additional correlations and possible bias. A more rigorous evaluation can be performed using pseudo-experiments or toy experiments. In the present

analysis, dedicated studies showed that the impact of the second-derivative regularisation on the uncertainty estimation is negligible for the chosen binning and signal-region definitions. Therefore, the standard uncertainty estimation is retained.

8.6.3 Reparametrisation and extraction of D_3

The PLU method described so far is used to extract the best-fit values of the normalisation factors $\hat{\vec{\mu}}$ to be applied to the nominal parton-level distribution. However, the primary goal of this measurement is the determination of the entanglement-sensitive observable D_3.

Recalling Equation 3.11 from Section 3.2.3,

$$D_3 = 3 \cdot \frac{\langle \cos\theta' \rangle}{\alpha_1 \alpha_2},$$

D_3 can be extracted from the mean value of the $\cos\theta'$ distribution. For a binned parton-level histogram, D_3 can be written in terms of the bin normalisation factors as:

$$D_3 = \frac{3}{\alpha_1 \alpha_2} \frac{\sum_{i=1}^{N_{\text{bins}}^{\text{parton}}} \mu_i \,\text{Centre}(i)\,\text{Yield}(i)}{\sum_{i=1}^{N_{\text{bins}}^{\text{parton}}} \mu_i \,\text{Yield}(i)}.$$

Here, $\text{Yield}(i)$ denotes the nominal yield in the i^{th} bin of the parton-level distribution $S^{\text{parton,nom}}$, $\text{Centre}(i)$ denotes the effective centre of the corresponding bin, and μ_i is the normalisation factor of the i^{th} bin of $S^{\text{parton,nom}}$, introduced previously. It is therefore evident that D_3 can be expressed directly in terms of the parameters describing the unfolded $\cos\theta'$ distribution at the parton level.

This relation can be rearranged such that the normalisation factor of one chosen parton-level bin r, can be expressed in terms of the remaining normalisation factors and D_3:

$$\mu_r = \frac{\sum\limits_{i=1, i\neq r}^{N_{\text{bins}}^{\text{parton}}} \text{Yield}(i)\,\mu_i \left(3\,\text{Centre}(i) - D_3\alpha_1\alpha_2\right)}{\text{Yield}(r)\left(D_3\alpha_1\alpha_2 - 3\,\text{Centre}(r)\right)}.$$

By replacing μ_r with this expression, the likelihood becomes a function of the remaining normalisation factors and of D_3 itself. Performing the unfolding fit with this reparametrisation therefore allows the direct extraction of D_3 as a fit parameter. Care must be taken to ensure that the extracted result does not depend significantly on the choice of the reparametrised bin r. This was explicitly tested in the present analysis, and the binning scheme, together with the regularisation choices, was optimised such that a stable and bin-independent determination of D_3 is obtained.

8.7 Systematic uncertainties

Various sources of systematic uncertainty are considered in this measurement. At the the time of writing, not all uncertainty sources had been fully validated within the fit framework. Therefore, only those uncertainties that were properly validated are included in the results presented in this thesis. Since this analysis is based on the software framework for Run-3 data analyses, the fast simulation

samples used primarily to derive the modelling systematic uncertainties are produced using the *AtlFast3* (AF3) framework. AF3 is the successor to AtlFast-II (AF2) [136], which was used in the threshold entanglement analysis. Compared with AF2, AF3 introduces significant improvements, in particular two parametric calorimeter simulation approaches that reduce simulation time while improving physics performance [256]. Therefore, the Run-2 MC event samples used in this analysis have been reprocessed using the Run-3 software framework.

The systematic uncertainties can be affected by statistical fluctuations arising from limited MC statistics in the simulated samples used to evaluate them. If propagated directly into the fit, these fluctuations could lead to unstable fit behaviour, induce spurious constraints on NPs, and introduce unphysical correlations between systematic uncertainties. To mitigate these effects, ***smoothing*** techniques are applied within TRexFitter to reduce sensitivity to statistical noise while preserving the underlying physical behaviour of the systematic variations. Different smoothing options exist depending on how aggressive and structured the smoothing needs to be for a given variation. In this analysis, smoothing is performed using the default *MaxVariation* algorithm [257]. The algorithm reduces statistical fluctuations by merging bins in regions of limited statistics and constrains the number of significant shape changes in the ratio of the systematic variation to the nominal prediction. For experimental uncertainties, this smoothing is applied such that bin-to-bin correlations among these uncertainties are maintained. This correlated smoothing approach avoids double-counting the underlying MC statistical fluctuations of the distributions, which are largely common across these uncertainties. In contrast, modelling uncertainties are smoothed independently in each bin, since they are derived from separate simulated samples defining the variations, and their associated MC statistical uncertainties must therefore be treated independently. For the PS and hadronisation uncertainty, however, smoothing was found to artificially induce shape distortions and enhance correlations with other uncertainties, resulting in an overly flat and inflated variation that caused the corresponding NP to become strongly constrained. As this behaviour is not physically motivated, the PS uncertainty is left unsmoothed.

Along with smoothing, ***symmetrisation*** is also applied to the systematic uncertainties. For one-sided systematic uncertainties, i.e. those for which only a single variation with respect to the nominal sample is available, the missing variation is constructed by reflecting the available variation about the nominal prediction. For systematic uncertainties with both upward and downward variations available, symmetrisation is performed by redefining the variations such that the upward and downward shifts are equal in magnitude and correspond to half the difference between the original up and down variations. These symmetrisation procedures are applied assuming that the upward and downward variations associated with a given uncertainty considered in this fit configuration are approximately symmetric.

8.7.1 Experimental

The main sources of experimental uncertainty considered in this analysis arise from the modelling of pile-up, electrons, muons, and jets. The treatment of these uncertainties follows that described in Section 7.9.1. All these uncertainties are applied on both signal and backgrounds.

For electrons, uncertainties related to the scale factors applied for reconstruction, identification, isolation, energy scale, and energy resolution are included. For muons, uncertainties are applied to the reconstruction, identification, isolation, and track-to-vertex association efficiences, accounting for both statistical and systematic components. Uncertainties in the momentum scale and resolution of CB muons are also included, together with contributions from the *sagitta* (track curvature) measurement.

Uncertainties related to the flavour-tagging performance of the GN2 algorithm are also taken into account. These uncertainties are propagated through variations of the flavour-tagging scale factors applied in simulation. Similar to the procedure used for the DL1r algorithm (Section 7.9.1), an eigenvector decomposition method is used to reduce the full covariance matrix of the calibration uncertainties into a set of independent variations. This results in 54, 23, and 41 eigenvector components for the b-, c-, and light-flavour jet scale factors, respectively. To account for the extrapolation of these calibration scale factors into the high-p_{T} regime, three additional uncertainty components are included for each jet flavour category.

The JES and JER uncertainties associated with the jets are defined in the same way as described in Section 7.9.1. Certain uncertainty components that were originally treated individually are combined into a single effective contribution. Additional uncertainty components are included to cover potential non-closure between the in-situ calibrations applied to full simulation and AF3 fast simulation samples. Further components account for uncertainties related to the jet-vertexing algorithm in the JES calibration and to calorimeter noise thresholds affecting both JES and JER. A single uncertainty component is also applied to the JVT discriminant.

8.7.2 $t\bar{t}$ modelling

Modelling uncertainties affecting the $t\bar{t}$ signal are also considered. These include uncertainties associated with the PS and hadronisation modelling, the recoil-to-top scheme, and the choice of the h_{damp} parameter. The definitions of these uncertainty sources are the same as those used in the threshold entanglement measurement (Section 7.9.2). In this analysis, they are evaluated by comparing the nominal prediction to alternative samples generated with varied settings. The main difference with respect to the threshold analysis is that the samples used here are produced with a newer software release. In addition to these uncertainties, several modelling uncertainties related to the variation of p_T^{hard} setting, PDFs, and renormalisation and factorisation scales are not yet included in the fit configuration. These are expected to be included and their impact tested prior to the finalisation of the analysis, as discussed in Section 8.11.

8.8 Event yields and differential distributions in the signal regions

The event yields at the reconstruction level for the different signal regions are summarised in Table 8.4. The expected $t\bar{t}$ signal dominates in all regions, contributing more than 90% of the total yield. Fewer events are observed in SR2 compared to SR1, reflecting the falling cross-section with increase in $m_{t\bar{t}}$. To account for this difference, the binning of the $\cos\theta'$ distributions within the fit setup is optimised separately in each region: SR1 is divided into six bins, while SR2 uses four bins. At the parton level, the distribution is evaluated in four bins.

Figure 8.3 shows the migration matrix, efficiency, and acceptance histograms used to construct the response matrix, also displayed in the figure, for SR1 e+jets. For the migration matrix, the content of each bin is stated explicitly. Each entry quantifies the probability for an event generated in a given parton-level bin to be reconstructed in a corresponding reconstruction-level bin. The probabilities sum to unity along each row or column. The largest values are concentrated along the diagonal, indicating that events are most likely to be reconstructed in the same bin as defined at the parton level.

The acceptance histogram shows that the fraction of reconstructed events that also satisfy the

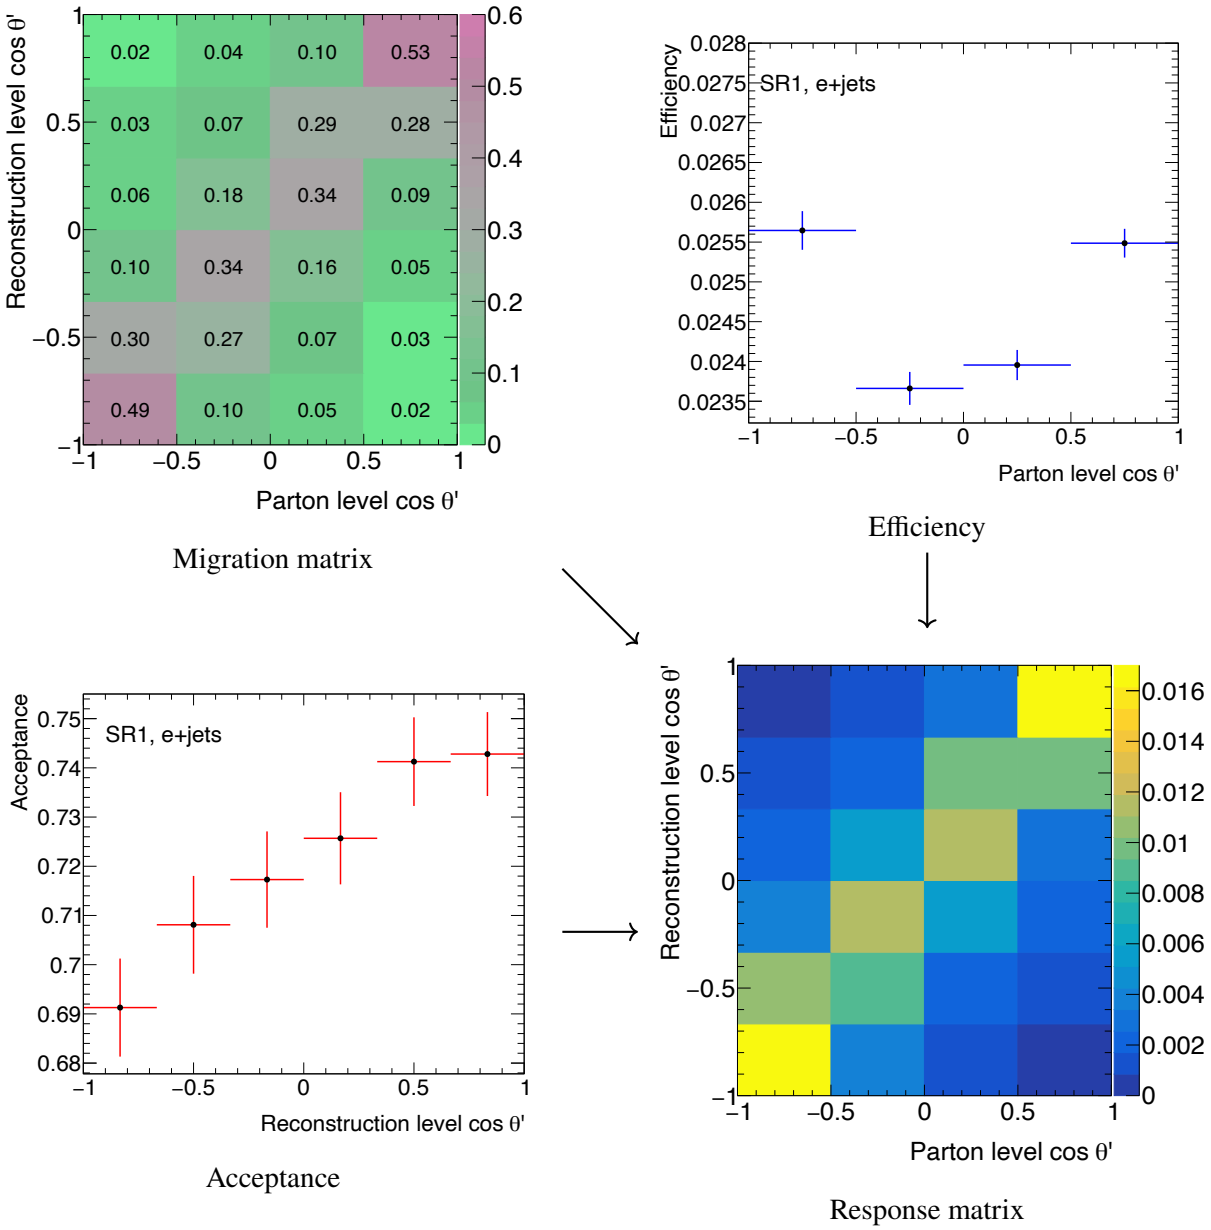


Figure 8.3: Migration matrix, acceptance, efficiency, and the resulting response matrix in SR1 e+jets. The z-axis of the migration matrix denotes the probability that an event reconstructed in a given reconstruction-level bin originates from a particular parton-level bin. These values are explicitly shown in the matrix itself. For the response matrix, the corresponding axis represents the same probability after accounting for detector acceptance and efficiency effects.

Table 8.4: Expected number of events per signal region prior to the fit, showing both the total expected yields and yields per process. The uncertainties include both statistical and systematic contributions.

	SR1		**SR2**	
	***e* + jets**	***μ* + jets**	***e* + jets**	***μ* + jets**
$t\bar{t}$	13000 ± 400	12000 ± 400	5820 ± 220	5190 ± 220
Single top	179 ± 12	167 ± 9	175 ± 7	161 ± 8
$t\bar{t}$+V/H	97.6 ± 3.4	87.6 ± 3.3	53.1 ± 1.7	46 ± 1.7
W+jets	22.7 ± 3.4	22.5 ± 2.7	33.9 ± 3.3	33 ± 4
Z+jets	4.5 ± 0.6	3.5 ± 0.4	6.3 ± 0.5	4.6 ± 0.4
Diboson	2.4 ± 0.4	1.9 ± 0.6	3.3 ± 0.4	2.7 ± 0.4
Total expectation	13300 ± 400	12300 ± 500	6090 ± 230	5440 ± 220

parton-level selection increases with $\cos\theta'$, reaching a maximum of approximately 74% in the highest bin. This behaviour could be attributed to the correlation between the angular variable and the kinematics of the $t\bar{t}$ system relevant for the boosted $m_{t\bar{t}}$ selection. Since $\cos\theta'$ encodes the relative orientation of the decay products in the top-quark rest frames, different values correspond to different event topologies, which affect how likely events are to remain within the fiducial phase space after reconstruction and binning. In particular, $\cos\theta' \approx -1$ corresponds to $\theta' \approx 180^\circ$, i.e. a back-to-back configuration of the decay products and hence the two spin analysers. In contrast, larger values of $\cos\theta'$ tend to correspond to configurations with aligned decay products. The latter configuration is found to have slightly higher acceptance than the former. This can be attributed to the dependence of the reconstruction performance on the detector acceptance. In regions of limited detector coverage, the reconstruction performance is reduced, leading to a smaller overlap between reconstruction- and parton-level selections, and thus to a reduced acceptance.

The efficiency histogram quantifies the fraction of events generated at the parton level that pass the reconstruction-level selection. The maximum efficiency is approximately 2.5%, indicating that only a small fraction of parton-level events are successfully reconstructed. This reflects the combined effects of reconstruction requirements, less than 100% b-tagging efficiency, and differences between the parton-level and reconstruction-level phase spaces. For example, in events with a high jet multiplicity, not all jets required to define the lepton+jets phase space may be reconstructed, leading to a loss of events at the reconstruction level. The reduced efficiency, together with the effects of acceptance and bin migration, results in a response matrix with bin contents significantly smaller than unity.

The corresponding histograms in SR2 e+jets region are shown in Figure 8.4. While the migration and acceptance behaviour remains identical to SR1, the efficiency is much lower, even though a coarser binning is used. This can be attributed to the more challenging reconstruction conditions in this region, due to increased overlap and reduced object separation. In addition, differences in the definition of the spin analysers at the two levels can also contribute. As a result, the overall response in this region is reduced compared to SR1, indicating that the unfolding procedure in SR2 is more statistically limited. Despite this, the chosen binning, together with second-order Tikhonov regularisation and the smoothing of systematic variations, ensures a stable unfolding procedure. Corresponding distributions for μ+jets regions show similar behaviour and are available in Appendix C.

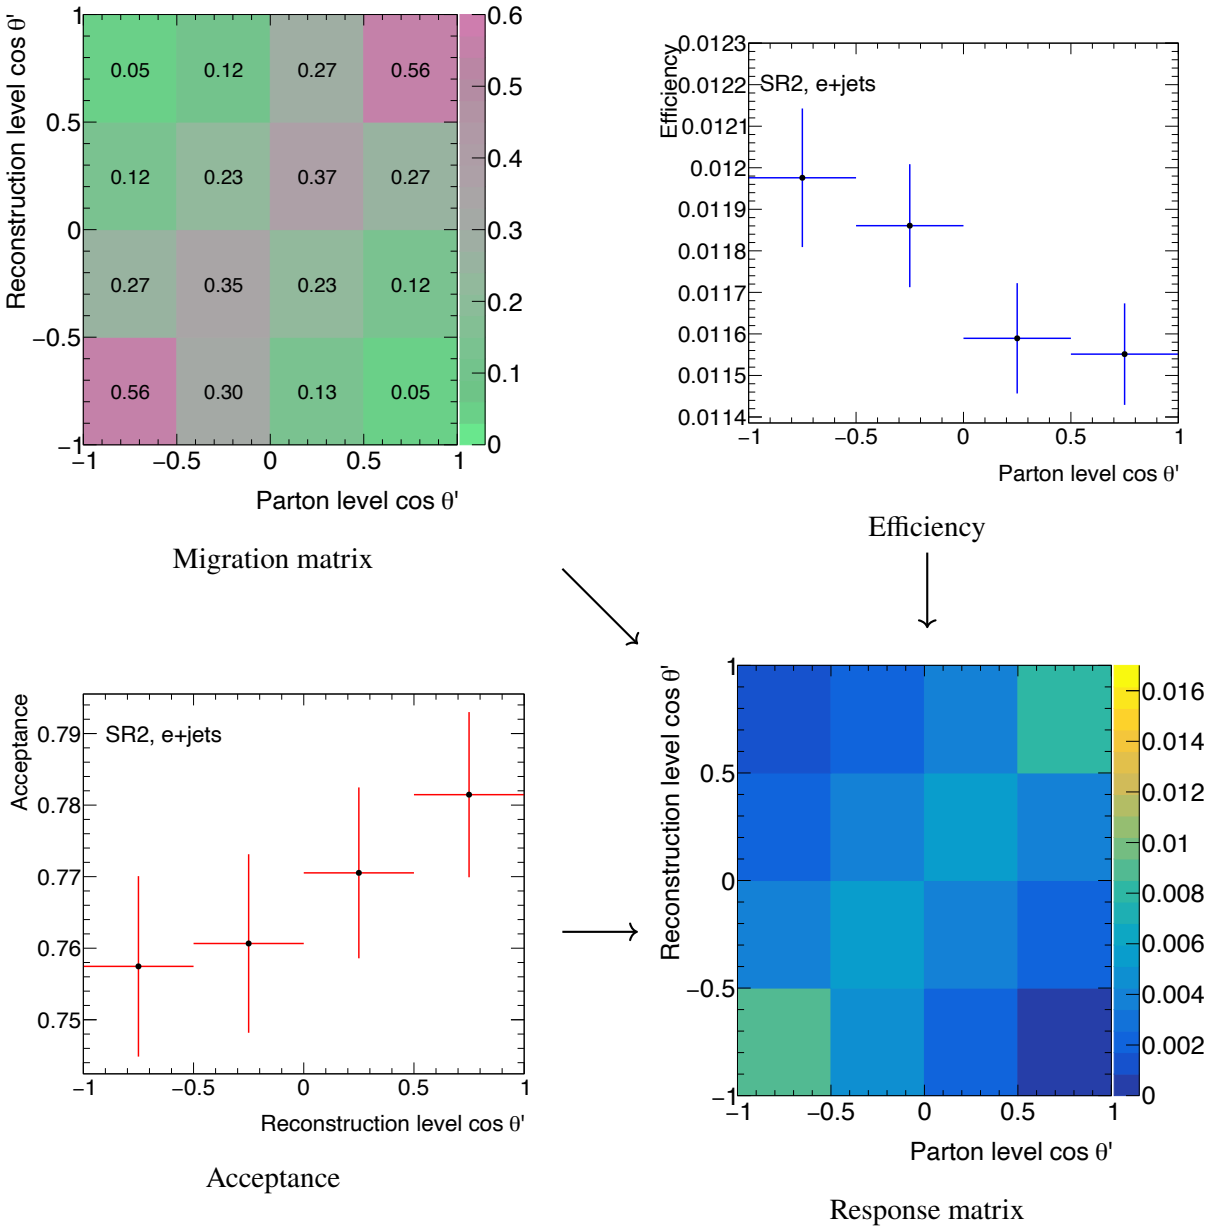


Figure 8.4: Migration matrix, acceptance, efficiency, and the resulting response matrix in SR2 e+jets. The z-axis of the migration matrix denotes the probability that an event reconstructed in a given reconstruction-level bin originates from a particular parton-level bin. These values are explicitly shown in the matrix itself. For the response matrix, the corresponding axis represents the same probability after accounting for detector acceptance and efficiency effects.

8.9 Results

The response matrices in each region are used to fold the nominal $t\bar{t}$ parton-level $\cos\theta'$ distribution to the reconstruction level in order to extract D_3 via the PLU procedure. Following the reparametrisation procedure explained in Section 8.6.3, in this fit the third parton-level bin is reparametrised to directly extract the value of D_3. This choice is motivated because compared to the edge bins of the $\cos\theta'$ distribution, the central bins exhibit larger bin-to-bin migration effects, as shown by the migration matrices in Figures 8.3 and 8.4. By reparametrising a central bin (i.e. the third parton-level bin), the more stable edge bins can be retained as independent degrees of freedom, while improving the overall stability of the unfolding description. This makes the reparametrisation of the third bin a natural choice in this case. As noted previously, given the binning scheme used in this analysis, the final results are independent of which bin is reparametrised.

Figures 8.5 and 8.6 show the pre- and post-fit reconstruction-level distributions in SR1 and SR2, respectively. As noted previously, all signal regions are blinded, meaning that the fit is not performed on real collision data. Instead, ***Asimov data*** is used. This dataset is constructed from the expected signal and background yields in each bin and provides a useful validation of the fit model prior to its application to real data. Therefore, no data points are shown in the figure.

The $t\bar{t}$ signal contribution in these figures corresponds to the folded distribution obtained by summing the contributions from all $t\bar{t}$ $\cos\theta'$ parton-level bins after applying the response matrix. The post-fit normalisation factors are then applied to the nominal parton-level distribution to extract D_3. It can be seen that for each signal region, the uncertainty band for the post-fit distribution is considerably smaller than the corresponding pre-fit distribution. This indicates that the fit is constraining several of the dominant NPs. A more detailed breakdown of which uncertainties are most strongly constrained, and the reasons for this behaviour, is given in Section 8.9.2.

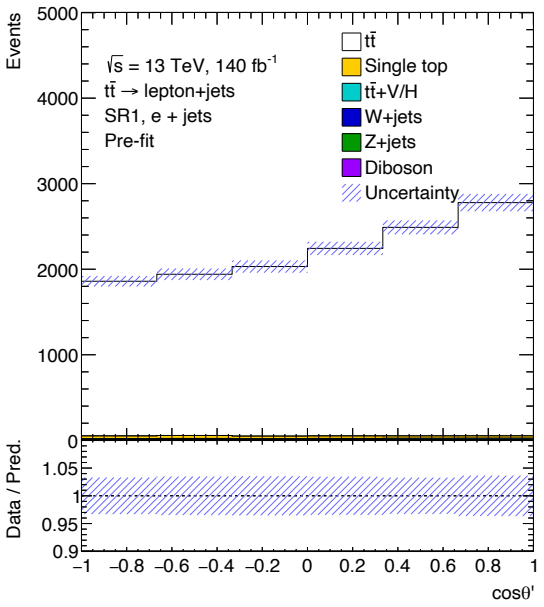
8.9.1 Expected value of D_3

Following the fit, the parton-level value of D_3 extracted via PLU is shown in Table 8.5. The post-fit values of the three normalisation factors, denoted by $\hat{\mu}_1$, $\hat{\mu}_2$, and $\hat{\mu}_4$, extracted as part of the fit, are also reported. Since the fit is performed on Asimov data, these parameters are expected to be unity, which is indeed observed.

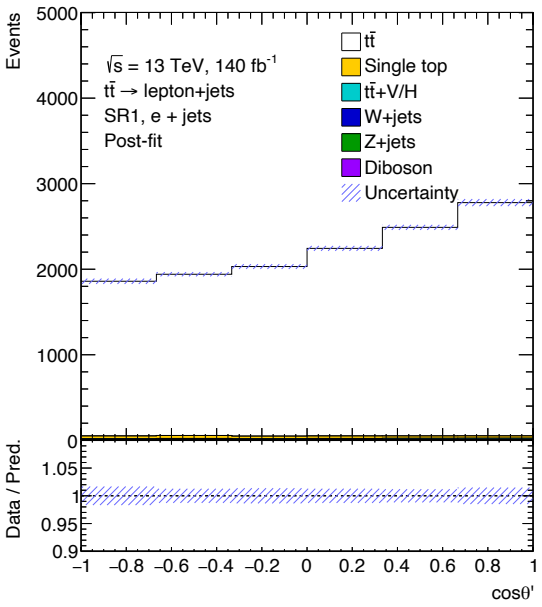
Table 8.5: Summary of the POIs extracted from the PLU fit. The quoted uncertainties are decomposed into statistical and systematic components, listed in that order.

POI	
D_3	-0.561 ± 0.020 ± 0.008
$\hat{\mu}_1$	$1.000\ ^{+0.028}_{-0.026} \pm 0.029$
$\hat{\mu}_2$	$1.000\ ^{+0.039}_{-0.038} \pm 0.032$
$\hat{\mu}_4$	$1.000\ ^{+0.020}_{-0.018} \pm 0.029$

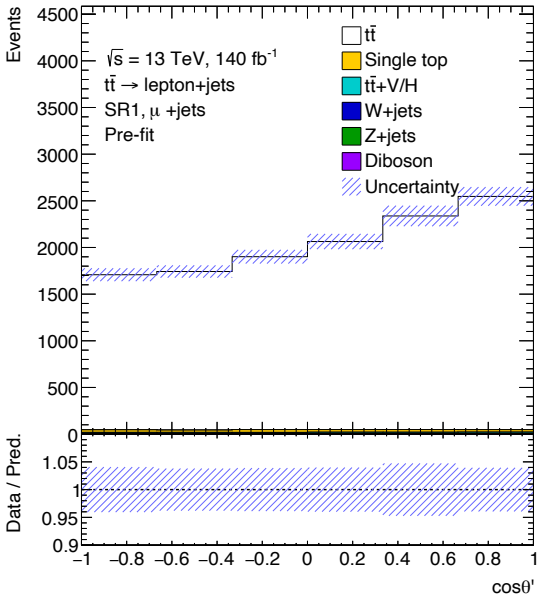
The extracted value of D_3 lies well below the entanglement limit of $-1/3$, indicating that the $t\bar{t}$ system is expected to be ***maximally entangled in the high $m_{t\bar{t}}$ regime***. This is consistent with the theoretical expectation that $t\bar{t}$ pairs are predominantly produced in a spin-triplet state in this boosted phase space (Section 3.2.3). Examining the uncertainties, the result is dominated by statistical uncertainties, as expected given the available event yields.



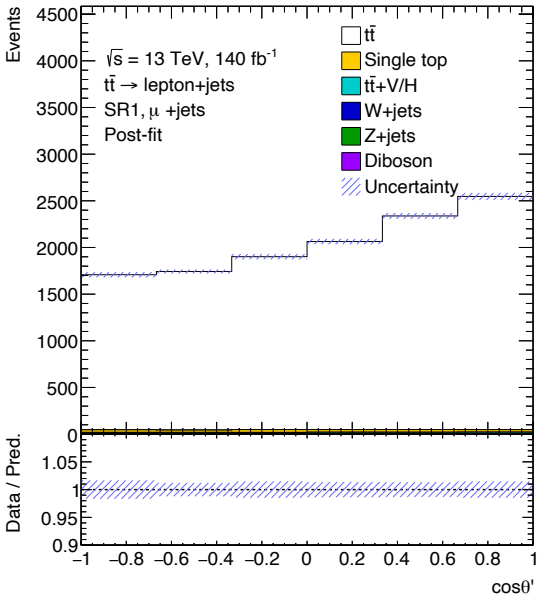
(a) Pre-fit distribution, SR1 e+jets



(b) Post-fit distribution, SR1 e+jets

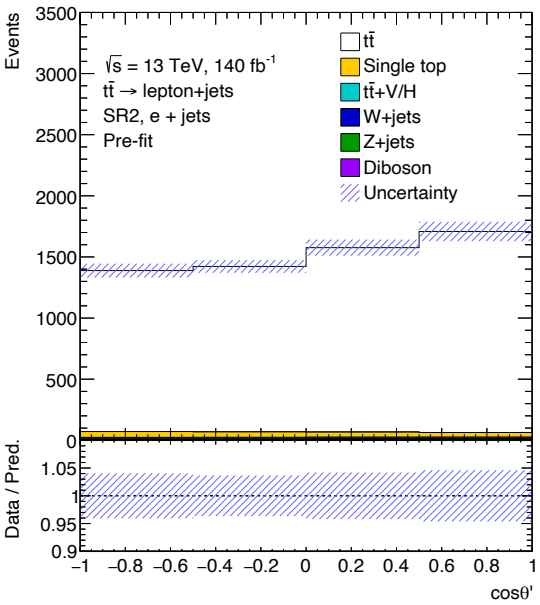


(c) Pre-fit distribution, SR1 μ+jets

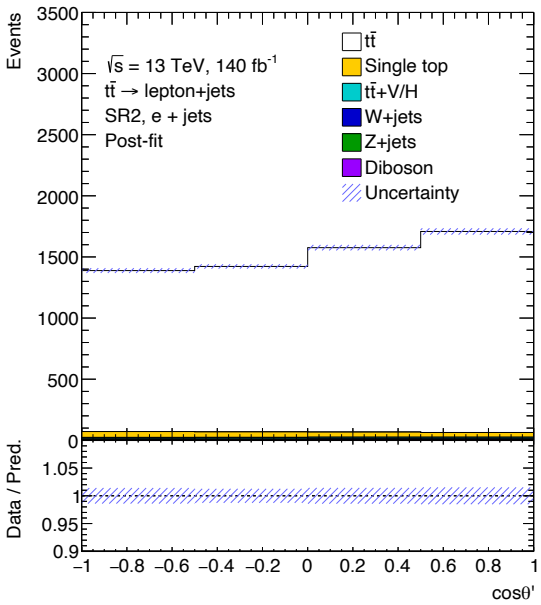


(d) Post-fit distribution, SR1 μ+jets

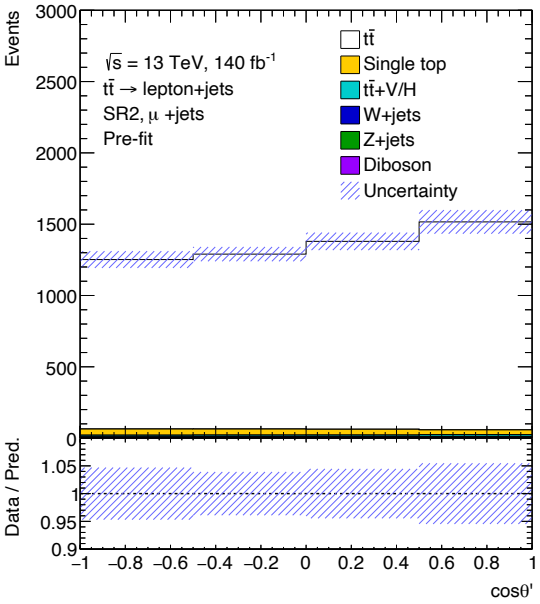
Figure 8.5: Pre- and post-fit distributions of $\cos\theta'$ in SR1, showing the MC signal and background contributions. The data used in the fit correspond to Asimov data. Hence, no data points are displayed. The lower panel shows the ratio of the data to the total MC prediction. The uncertainty band includes the full set of statistical and systematic uncertainties considered in the fit.



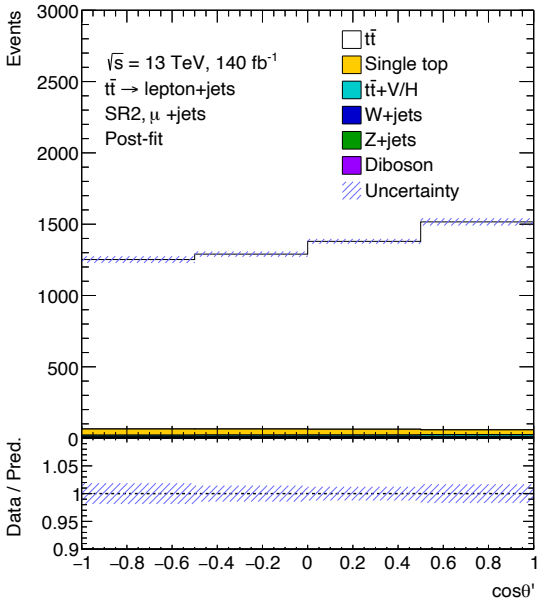
(a) Pre-fit distribution, SR2 e+jets



(b) Post-fit distribution, SR2 e+jets



(c) Pre-fit distribution, SR2 μ+jets



(d) Post-fit distribution, SR2 μ+jets

Figure 8.6: Pre- and post-fit distributions of $\cos\theta'$ in SR2, showing the MC signal and background contributions. The data used in the fit correspond to Asimov data. Hence, no data points are displayed. The lower panel shows the ratio of the data to the total MC prediction. The uncertainty band includes the full set of statistical and systematic uncertainties considered in the fit.

Since the measurement is performed using two orthogonal signal regions, the corresponding results in SR1 only and SR2 only are shown in Table 8.6 for comparison, alongside the combined fit result. SR2 carries the largest uncertainty on D_3, as expected due to its lower efficiency and overall response. When combined with SR1, it provides slight reduction in the uncertainty of the combined fit result relative to SR1 alone. This demonstrates that the combination benefits from simultaneously capturing both two-jet and three-jet topologies, thereby improving the overall statistical precision. The relative benefit of the combination is expected to become more pronounced in more boosted regions defined by tighter $m_{t\bar{t}}$ selections. In such regions, the statistical limitation can be partially mitigated by exploiting the complementary sensitivity of both two-jet and three-jet topologies. Furthermore, combining this result with Run-3 data is expected to not only improve the statistical precision, but also to enhance the complementarity in event topology and systematic uncertainties.

Table 8.6: Summary of the POIs extracted from the PLU fit for the different signal regions and their combination. The quoted uncertainties are decomposed into statistical and systematic components, listed in that order.

POI	SR1	SR2	Combination
D_3	-0.561 ± 0.022 ± 0.009	-0.561 ± 0.039 ± 0.039	-0.561 ± 0.020 ± 0.008
$\hat{\mu}_1$	$1.000\ ^{+0.029}_{-0.026} \pm 0.032$	$1.000\ ^{+0.100\ +0.073}_{-0.124\ -0.066}$	$1.000\ ^{+0.028}_{-0.026} \pm 0.029$
$\hat{\mu}_2$	$1.000\ ^{+0.041}_{-0.038} \pm 0.039$	$1.000\ ^{+0.170}_{-0.158} \pm 0.043$	$1.000\ ^{+0.039}_{-0.038} \pm 0.032$
$\hat{\mu}_4$	$1.000\ ^{+0.020}_{-0.017} \pm 0.032$	$1.000\ ^{+0.064\ +0.062}_{-0.071\ -0.060}$	$1.000\ ^{+0.020}_{-0.018} \pm 0.029$

8.9.2 Impact of systematic uncertainties in the fit

As introduced in Section 8.6.1, pulls are defined in the profile-likelihood fit and indicate the deviation of a NP from its nominal value, together with its post-fit constraint. In the case of a fit to Asimov data, the nominal expectation is that the pulls are centred at zero, since the data is constructed to exactly match the prediction of the model. While the pulls themselves are therefore expected to be trivial, the post-fit constraints of the NPs provide valuable information about the sensitivity of the fit to different systematic uncertainties and can give insight into what may be observed when performing the fit with real collision data. In the following, results from the combined fit are discussed.

Figure 8.7 shows the pull plots of the $t\bar{t}$ modelling systematics implemented so far. A moderate level of constraint is observed across several parameters, with the largest constraint seen in the PS and hadronisation uncertainty. This is not unusual, since the uncertainty is known to be difficult to model accurately, as also seen in the threshold entanglement measurement (Chapter 7). The constraint indicates that the chosen parametrisation induces sizeable deviations in certain regions of phase space, suggesting that an improved modelling of this uncertainty could be beneficial. Figures 8.8 to 8.13 show the pull plots for the remaining systematic uncertainties. For the muons (Figure 8.11), the systematic uncertainty components of the reconstruction, isolation, and TTVA efficiencies are shown separately, corresponding to a decomposition of the full muon efficiency uncertainty into its individual contributions. This breakdown is used to study the sensitivity of the analysis to the individual components as well as their potential correlations and profiling effects. For the future iterations of this analysis, the breakdown would potentially be reduced to a simpler scheme by summing small uncertainties. No notable constraints are observed for most parameters, with the exception of a slight constraint on the muon isolation efficiency in the $m_{\ell\ell}$ window (Figure 8.11). This uncertainty is

related to the variation of the invariant-mass window of the $\mu^+\mu^-$ system used in the derivation of muon isolation efficiency scale factors from $Z \to \mu^+\mu^-$ and $J/\psi \to \mu^+\mu^-$ events [203].

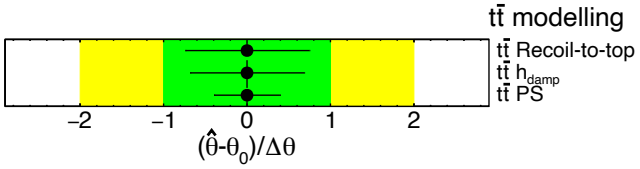


Figure 8.7: Pulls and constraints for NPs associated with $t\bar{t}$ modelling uncertainties. Black points indicate pulls, while horizontal bars represent constraints. The green and yellow bands correspond to the 1σ and 2σ ranges.

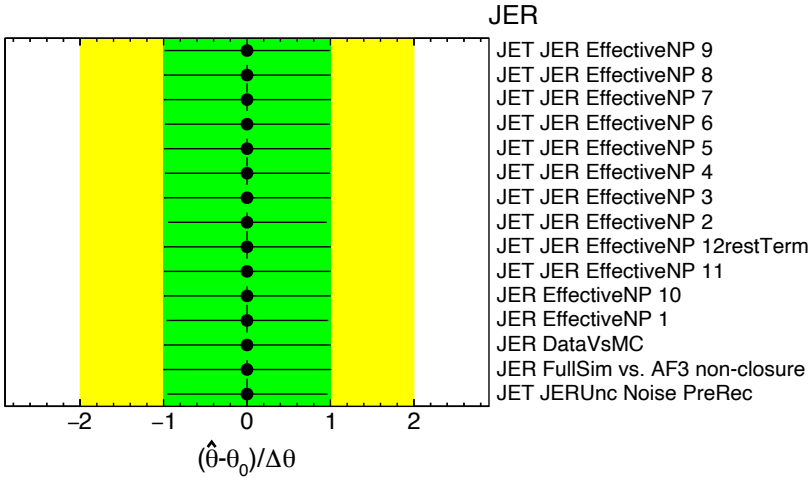


Figure 8.8: Pulls and constraints for NPs corresponding to the JER systematic uncertainties.

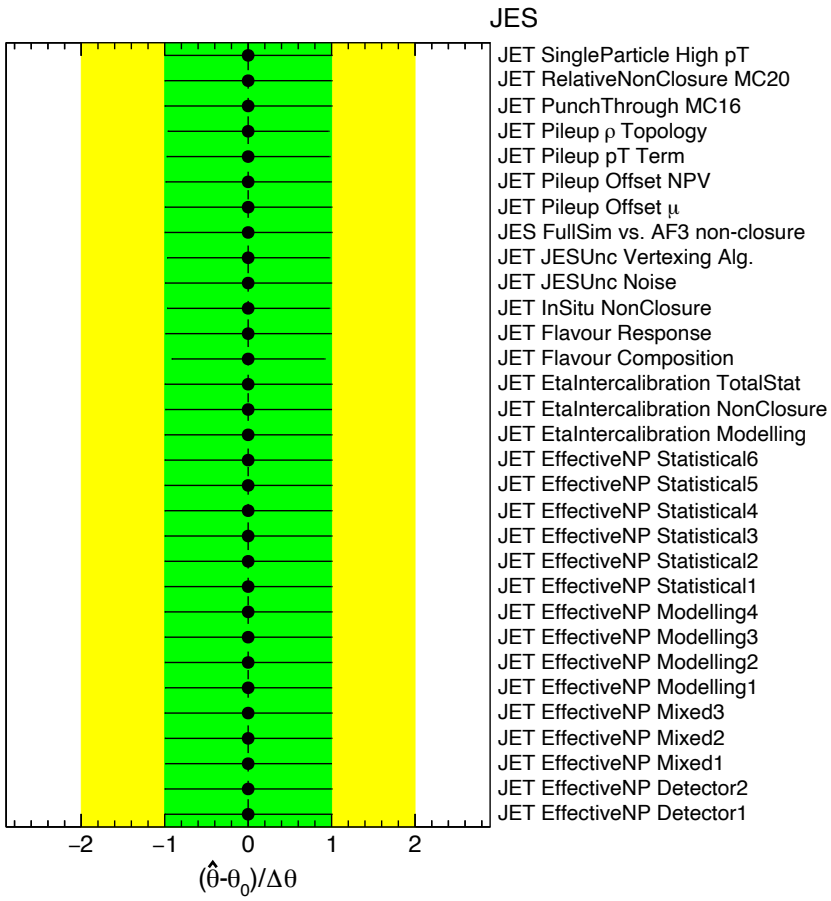


Figure 8.9: Pulls and constraints for NPs corresponding to the JES systematic uncertainties.

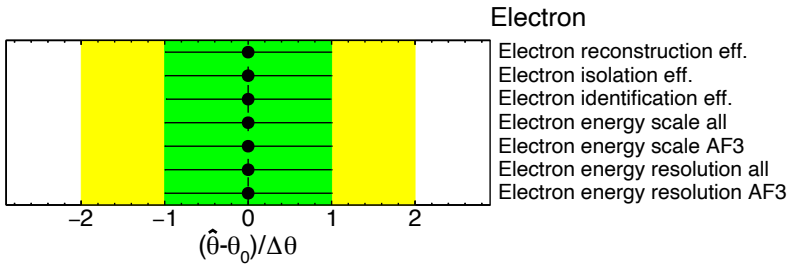


Figure 8.10: Pulls and constraints for NPs corresponding to the systematic uncertainties associated with electrons.

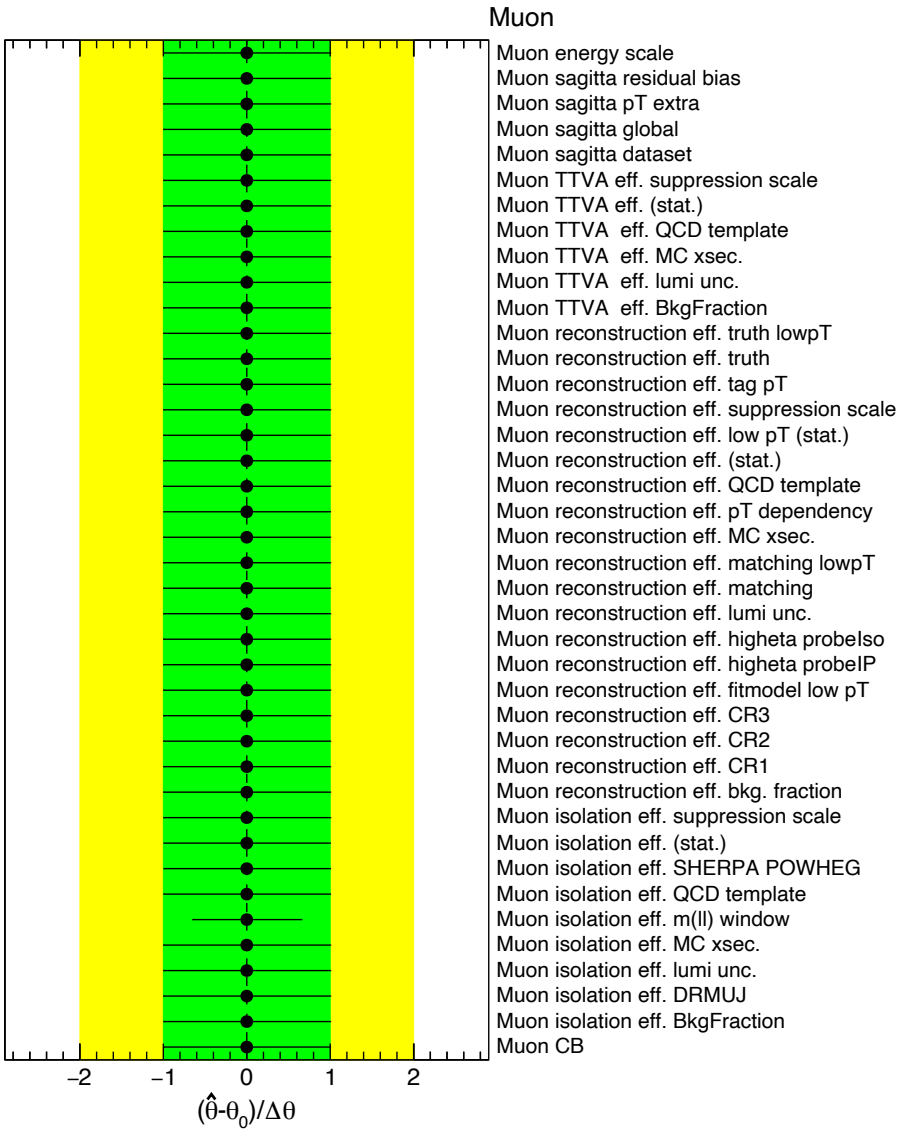


Figure 8.11: Pulls and constraints for NPs corresponding to the systematic uncertainties associated with muons. The systematic components of the reconstruction, isolation, and TTVA efficiencies are shown separately, as explained in the text.

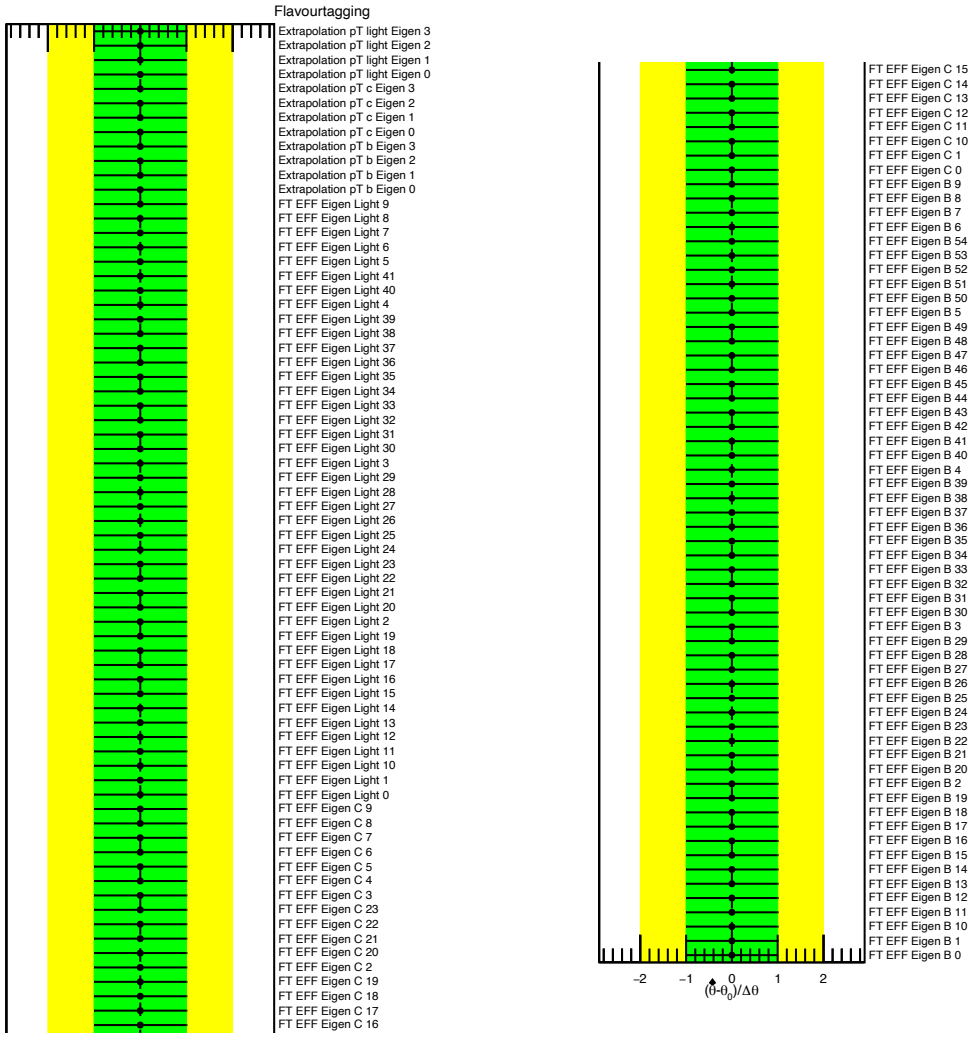


Figure 8.12: Pulls and constraints for NPs corresponding to the flavour-tagging systematic uncertainties. FT EFF denotes flavour tagging efficiency uncertainties. The components labelled Eigen B, Eigen C, and Eigen Light correspond to the eigenvector variations associated with b-, c-, and light-flavour jets, respectively. The corresponding high-p_T extrapolation uncertainties are denoted by pT b Eigen, pT c Eigen, and pT light Eigen.

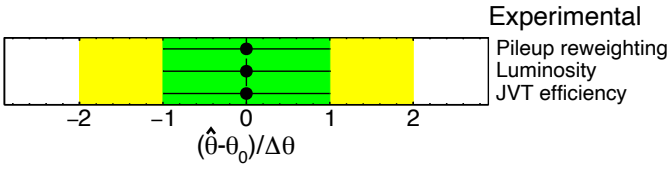


Figure 8.13: Pulls and constraints for NPs of experimental uncertainties related to pile-up reweighting, luminosity, and JVT efficiency.

To further quantify the influence of systematic uncertainties on the parameter of interest, D_3, their impact is evaluated. The definition of impact can vary [258], and in this analysis it is derived from the covariance matrix obtained from the fit under the Gaussian approximation [259]. Recalling Equation 8.3, the covariance matrix is given by the inverse of the Hessian of the log-likelihood at its maximum. Each diagonal element of the covariance matrix represents the variance of a given fit parameter, from which its uncertainty can be estimated (Equation 8.4). Each off-diagonal element, on the other hand, quantifies how two parameters vary together in a fit, $\hat{V}_{ij} = \mathrm{Cov}(\hat{\alpha}_i, \hat{\alpha}_j)$. By normalising the covariance matrix element-wise, the corresponding ***correlation matrix*** is obtained, with elements given by $\rho_{ij} = \frac{\hat{V}_{ij}}{\sqrt{\hat{V}_{ii}\hat{V}_{jj}}}$. The correlation matrix therefore quantifies the correlations between pairs of fit parameters, with values ranging from -1 to 1. Figure 8.14 shows the correlation matrix, illustrating the correlations (in percentage) between the POIs and NPs included in the fit. Among the strongest correlations involving D_3, a moderate anti-correlation with the PS uncertainty of approximately -15.4% is observed, together with a correlation of 16.4% with the recoil-to-top uncertainty.

The impact of a NP on a POI can then be derived from the post-fit covariance matrix through the corresponding correlations and parameter uncertainties [259]. For a given POI a, the impact I of a NP n is estimated as:

$$I = \sigma^a_{\text{total}} \rho_{a,n} \sigma^n,$$

where σ^a_{total} is the total uncertainty on the POI, $\rho_{a,n}$ is the correlation between the POI and the NP, and σ^n is the post-fit uncertainty of the NP. The ranking plots showing the post-fit impacts of the leading systematic uncertainties on D_3 and the other normalisation factors are summarised in Figure 8.15.

The JET Flavour Response is found to have the largest impact on D_3, as seen in the ranking plot (Figure 8.15(a)). This behaviour can be understood by examining its effect on the reconstructed $\cos\theta'$ distribution. Figures 8.16 and 8.17 show the folded $t\bar{t}$ signal distributions for this systematic across the four parton-level bins in the μ+jets region of SR1 and SR2, respectively (similar distributions are also seen for e+jets). The up and down variations of this systematic are shown relative to the nominal prediction. In most cases, the variations lead to approximately uniform shifts across bins. However, for specific configurations, particularly for the folded distributions of the second and fourth parton-level bins in SR1, they induce non-uniform, bin-dependent changes. The deviation is generally largest in the fourth bin of the folded distributions. These deviations do not correspond to a simple reweighting of a single template, but instead introduce a change in the overall shape of the reconstructed distribution. Since D_3 is extracted from the distribution corresponding to the sum of all these folded distributions, such shape changes can significantly affect its determination, leading to a large impact. The comparatively larger deviation in the fourth bin for most of the folded distributions could also explain why the JET Flavour Response uncertainty is one of the highly ranking uncertainties for the corresponding normalisation factor μ_4 (Figure 8.15(d)). A similar procedure can be followed to understand the origin of the large impact of the other leading systematic uncertainties. Overall, the dominant impacts arise from jet and flavour-tagging systematic uncertainties, with few $t\bar{t}$ modelling uncertainties also appearing among the leading contributors. This behaviour is expected given the boosted topology, with the final state dominated by jets together with the challenge of correctly identifying the flavour of a given jet.

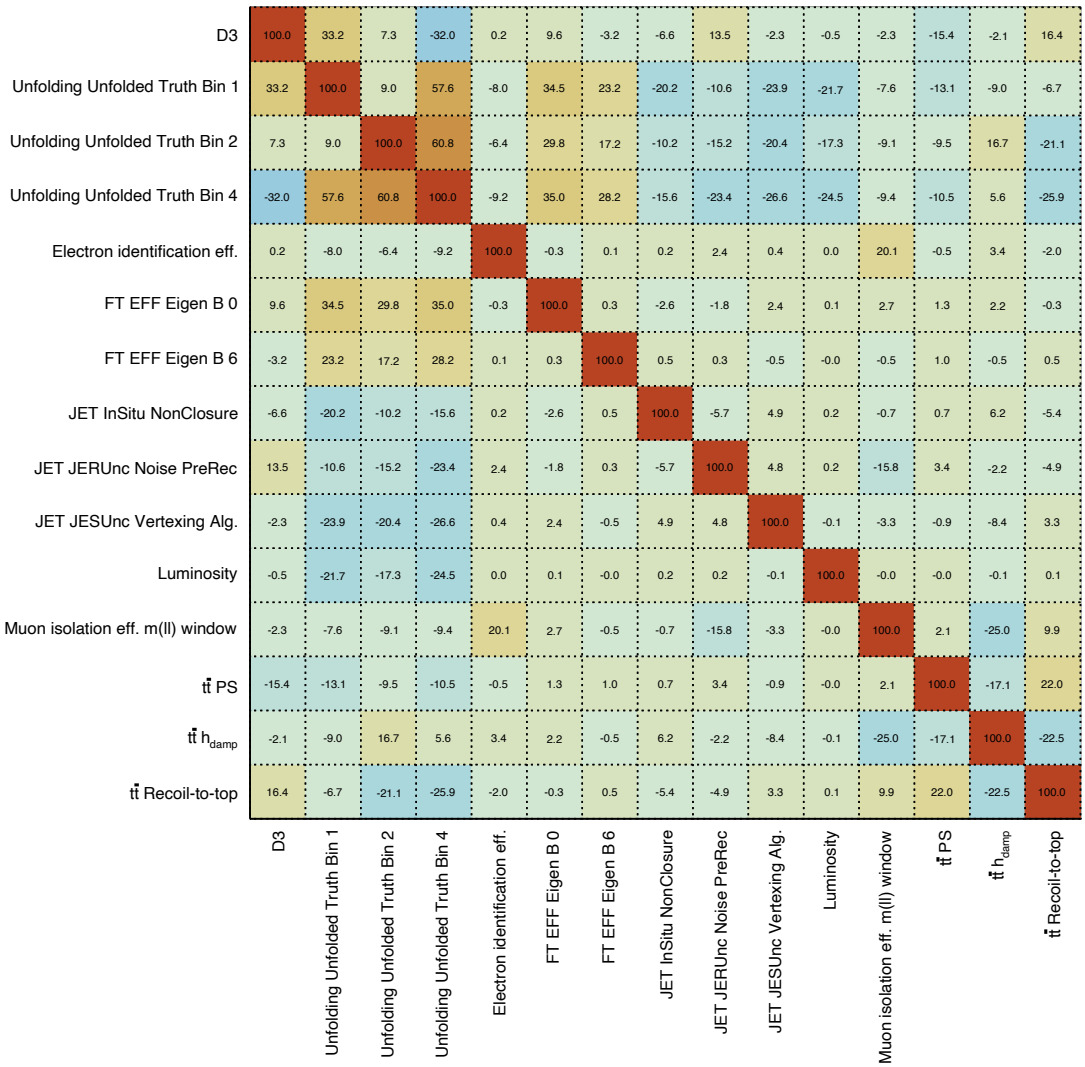
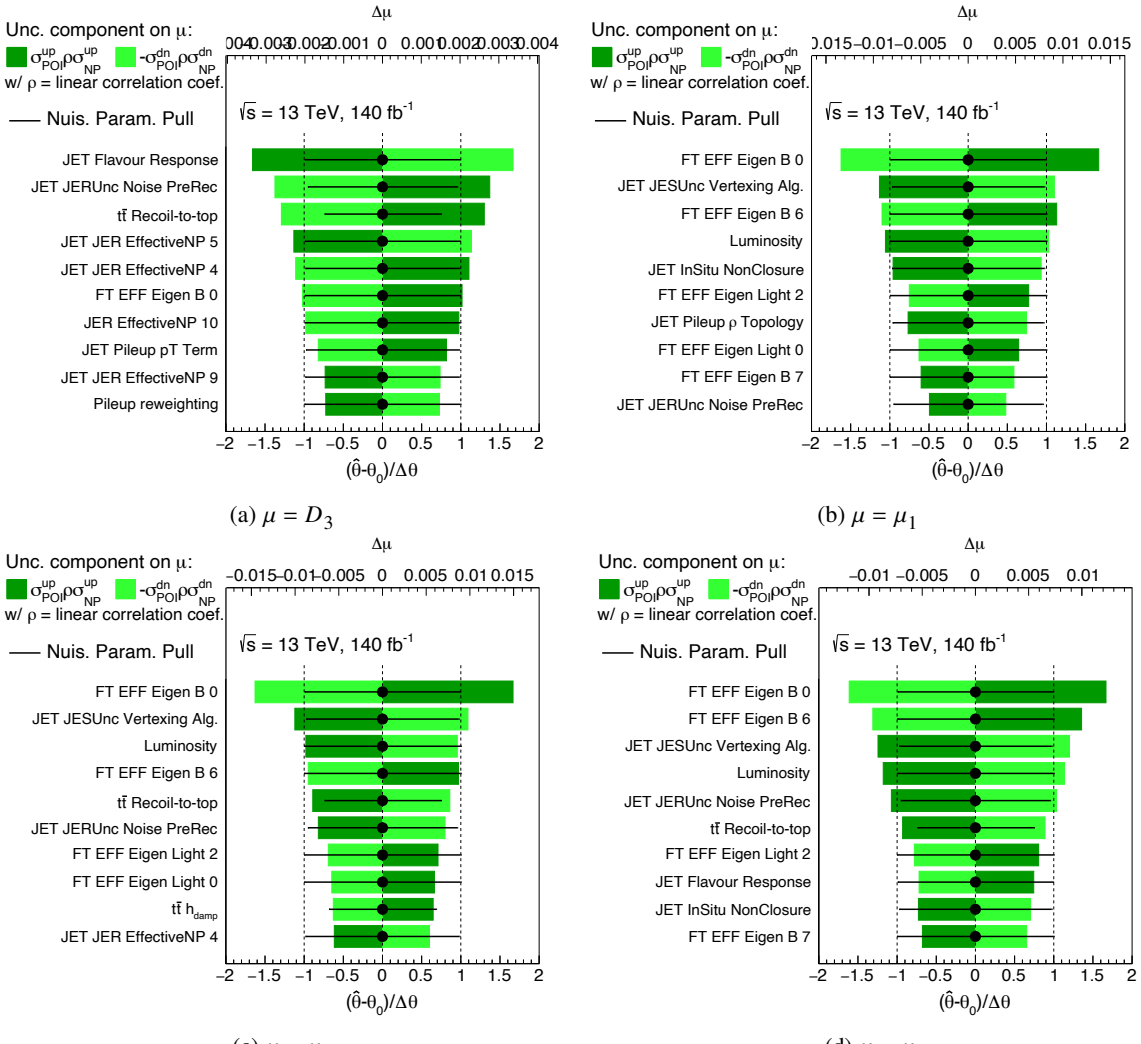


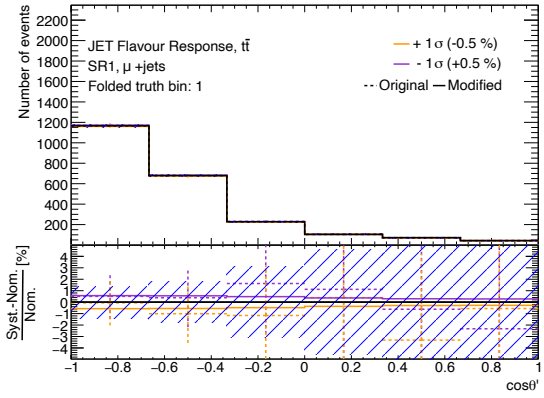
Figure 8.14: Correlation matrix for the POIs and NPs in the fit. The values represent the correlation coefficients between pairs of parameters, expressed in percentage. Only parameters with at least one correlation exceeding 20% in magnitude are shown.



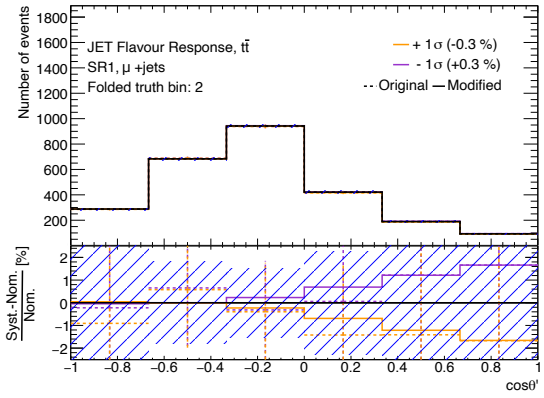
(a) $\mu = D_3$ (b) $\mu = \mu_1$

(c) $\mu = \mu_2$ (d) $\mu = \mu_4$

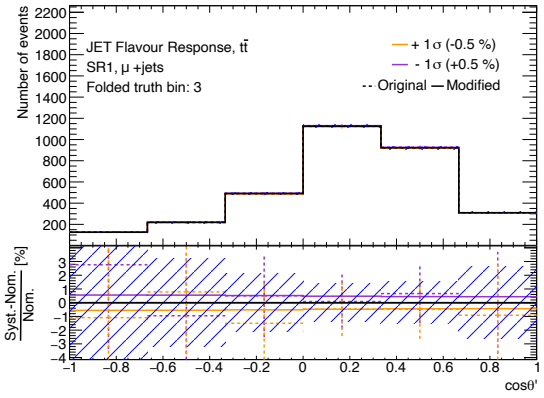
Figure 8.15: Post-fit ranking plot showing the impact of the 15 most significant uncertainty contributions to the extracted value of the POI, μ. The post-fit up ($\sigma^{\text{up}}_{\text{POI}}\,\rho\,\sigma^{\text{up}}_{\text{NP}}$) and down ($\sigma^{\text{down}}_{\text{POI}}\,\rho\,\sigma^{\text{down}}_{\text{NP}}$) impacts are shown. Pulls and constraints on the corresponding NPs are also indicated (black points with horizontal bars), with the parameter axis shown at the bottom.



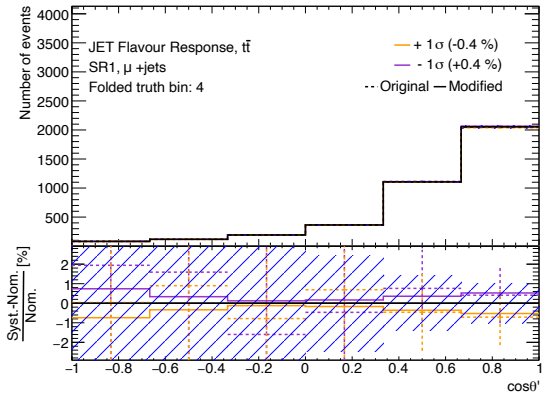
(a) Parton level bin 1



(b) Parton level bin 2

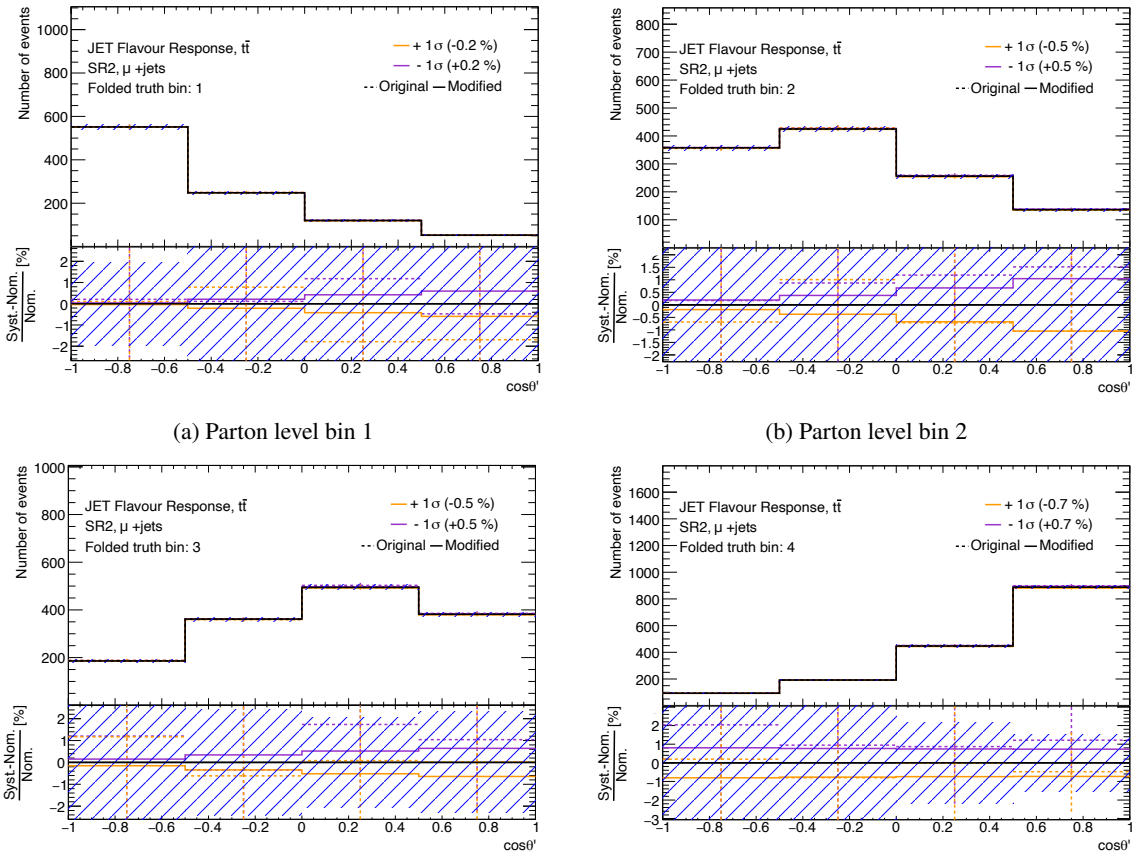


(c) Parton level bin 3



(d) Parton level bin 4

Figure 8.16: Shape of the JET Flavour Response systematic for a given parton level bin of the $\cos\theta'$ distribution of the $t\bar{t}$ signal folded to the reconstruction level in SR1, μ +jets channel. The up variation is given by the orange lines ($+1\sigma$) and the down variation is given by the violet lines (-1σ). The dashed lines correspond to the raw systematic distributions and the solid lines represent the smoothed and symmetrised distributions.



(a) Parton level bin 1

(b) Parton level bin 2

(c) Parton level bin 3

(d) Parton level bin 4

Figure 8.17: Shape of the JET Flavour Response systematic for a given parton level bin of the $\cos\theta'$ distribution of the $t\bar{t}$ signal folded to the reconstruction level in SR2, μ +jets channel. The up variation is given by the orange lines ($+1\sigma$) and the down variation is given by the violet lines (-1σ). The dashed lines correspond to the raw systematic distributions and the solid lines represent the smoothed and symmetrised distributions.

8.10 Comparison with entanglement measurement by CMS

As discussed in Section 3.2.3, the first measurement of entanglement in the boosted regime at the LHC was performed by CMS in the lepton+jets final state [24]. In contrast to the measurement presented in this thesis, in the CMS measurement the down-type jet is used as the hadronic spin analyser. The corresponding results are shown in Figure 8.18. In this context, D_3 is denoted as $\tilde{D}$ and is defined in the helicity basis, such that $\tilde{D} > 1/3$ implies that the $t\bar{t}$ system is entangled. This entanglement limit of $\tilde{D} = 1/3$ corresponds to the boundary separating the green band, which represents separable states, from the rest of the $\tilde{D}$ phase space. Entanglement was observed with a significance exceeding 5σ in a boosted phase space (left panel of the figure) defined similarly to that used in this thesis, with the observed value of $\tilde{D}$ reported as $\tilde{D} = 0.652 \pm 0.052$. The CMS results indicate larger degree of entanglement than the preliminary results obtained in the context of this thesis. This difference is not unexpected, since the CMS analysis employs a hadronic spin analyser that is more sensitive to the spin information of the hadronically decaying top quark than either the W boson or the optimal hadronic polarimeter used in this thesis. The uncertainties in the CMS result are largely dominated by the statistical contribution, with the leading systematic uncertainties arising from effects related to the top-quark mass, renormalisation and factorisation scales, jet energy scale, and b-tagging efficiency [24].

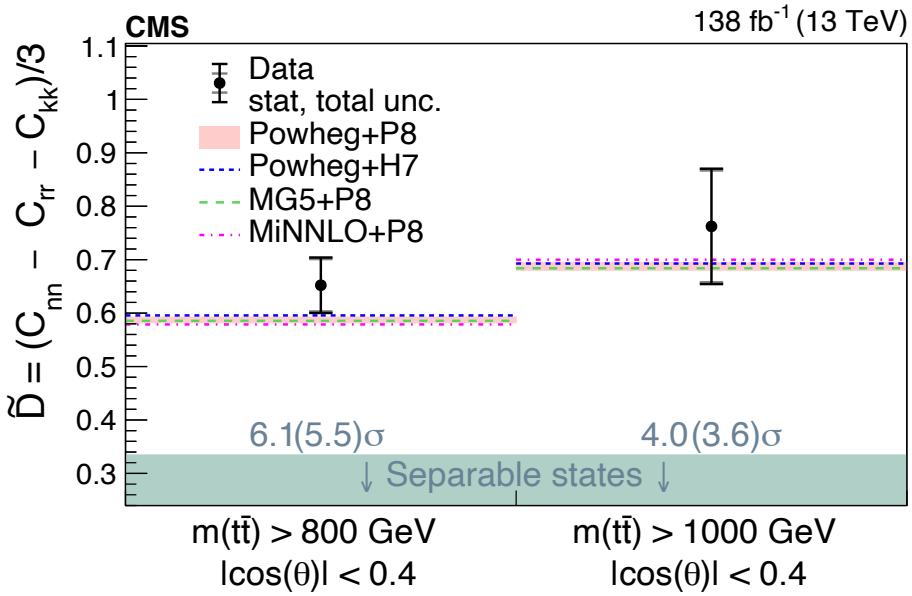


Figure 8.18: Measurement of $\tilde{D}$ (corresponding to D_3 in this thesis) by CMS in the lepton+jets boosted regime. The points show statistical (inner) and total (outer) uncertainties, compared to predictions from Powheg+Pythia, Powheg+Herwig, MadGraph5_aMC@NLO+Pythia, and MiNNLO+Pythia. The observed (expected) significance relative to separable states (green band) is quoted in units of σ. The boundary of the green band corresponds to the entanglement threshold $\tilde{D} = 1/3$, above which the $t\bar{t}$ system is entangled [24].

8.11 Outlook

With the baseline unfolding strategy for extracting D_3 already established, and the current results indicating sensitivity to entanglement in the boosted regime defined by $m_{t\bar{t}} > 800\,\text{GeV}$, this analysis is steadily progressing towards an application to real data. Before looking at data, however, several aspects of the analysis remain under active development.

The unfolding strategy presented in this thesis currently employs the W boson and the optimal hadronic polarimeter as spin analysers of the hadronically decaying top quark. However, it is known that the down-type quark has the highest spin analysing power and could therefore improve the sensitivity to entanglement if incorporated into the unfolding procedure. To achieve this, ongoing studies are investigating the use of c-tagging information from the GN2 algorithm to define a detector-level region enriched in correctly identified down-type jets. With the tagging capabilities of GN2, direct c-tagging can now be applied with efficiencies reaching approximately 50%. This allows a constituent of the RC jet associated with the hadronically decaying top quark to be identified as originating from a charm quark and use it in the down-type jet assignment.

The down-type jet definition can therefore be implemented within the fit model by introducing an additional category within SR1 for events in which the RC jet contains three constituents. In such events, if one constituent jet is identified as b-tagged and another as c-tagged, the remaining jet is assigned as the down-type jet candidate. This category can ultimately be unfolded to a corresponding parton-level definition of the down-type quark. Events that do not have a c-tagged jet continue to use the optimal hadronic polarimeter, ensuring orthogonality between the signal regions such that all categories can be combined consistently in the fit. This configuration, including the down-type jet definition, is planned to be studied further within the timeline of the analysis. If it results in an improved sensitivity that balances potential additional systematic uncertainties coming from the tagging itself, it may be adopted as one of the final signal-region definitions used in the measurement. An important aspect of these studies though lies in validating whether the current c-tagging calibration remains applicable in the boosted topology relevant for this analysis. Although calibrated c-tagging WPs are available, the corresponding flavour-tagging scale factors that are applied to MC simulations are derived primarily in resolved topologies, where the jets originating from the hadronically decaying top quark are well separated. It therefore remains to be established whether these calibrations remain valid in the boosted regime. If this is not the case, a dedicated calibration of the c-tagger for boosted topologies would need to be developed before such a signal-region with the down-type definition can be fully implemented.

The analysis currently does not impose requirements on the missing transverse momentum, E_T^{miss}, or on the transverse mass of the W boson, m_T^W. Introducing lower-bound requirements on these variables can help suppress backgrounds arising from non-prompt or fake leptons, while also ensuring consistency with the presence of a neutrino from a leptonic W-boson decay. In particular, a minimum requirement on m_T^W helps suppress events in which the reconstructed $\vec{E}_{\text{T}}^{\text{miss}}$ does not originate from a leptonic W-boson decay. The optimisation of these selections is currently under investigation and depends strongly on the strategy used to estimate non-prompt and fake-lepton backgrounds. As in the threshold entanglement measurement, these background contributions are estimated using a data-driven technique. However, in the boosted regime the selections on E_T^{miss} and m_T^W must be carefully chosen to retain sufficient statistical power in an already statistically limited phase space. Once this background estimation is finalised, optimised selections on the two variables are expected to improve the purity of the signal regions. Although the modelling of this background is challenging,

its overall contribution is expected to be small, and its inclusion is not expected to significantly alter the current results.

Beyond these developments, a major remaining task is the incorporation of the full set of systematic uncertainties expected to affect the measurement. The most important among these are expected to be the $t\bar{t}$ modelling uncertainties. While some of them have already been included in the current fit configuration, with PS and hadronisation being among the dominant contributions, several additional effects remain to be studied. These include variations of the p_T^{hard} setting, top-quark mass, colour reconnection, PDF uncertainties, FSR, ISR, and renormalisation and factorisation scale variations. Based on the behaviour observed in the current fit configuration and on previous results in similar topologies [24], these uncertainties are expected to not have a substantial impact on the sensitivity relative to the current result. Furthermore, the inclusion of the full systematic uncertainty model, through the introduction of additional NPs, is expected to increase the flexibility of the fit. This increased freedom is also expected to result in weaker constraints on individual NPs. Nevertheless, the precise impact of these uncertainties remains to be evaluated. Overall, the measurement result is expected to remain statistically dominated.

CHAPTER 9

Conclusion and outlook

"Everything is possible. The impossible takes longer."

— Dan Brown, Digital Fortress

The body of knowledge that shapes our understanding of the world around us is continuously evolving, and the work presented in this thesis contributes in a small but meaningful way to this progress. While the top quark is a well-established component of the SM, it also provides a unique opportunity to probe quantum phenomena in a high-energy environment. In particular, the study of $t\bar{t}$ quantum entanglement offers a novel perspective on top quarks, placing them within the broader context of quantum information science. The work presented in this thesis contributes to this growing body of research that bridges high-energy physics and quantum information science.

The threshold entanglement measurement marks the first observation of entanglement at the $t\bar{t}$ threshold in the lepton+jets channel by ATLAS (Chapter 7). Ongoing efforts in the boosted regime (Chapter 8) show strong potential for a first observation by ATLAS in that phase space as well. Together, these measurements highlight the rich structure of the spin configuration of $t\bar{t}$ pairs, which varies across kinematic regimes, and provide new opportunities to scrutinise the SM. Beyond validating previous results in this sector [17, 23, 24], these studies stand out through the use of more sophisticated $t\bar{t}$ modelling and improved analysis strategies. For example, the threshold entanglement measurement is the first entanglement analysis at the LHC to incorporate an advanced treatment of $t\bar{t}$ quasi-bound-state effects [95], while the boosted analysis explores alternative choices of hadronic spin analysers, thereby broadening the methodological landscape. Both measurements focus on the lepton+jets channel, which offers a favourable balance between statistical power and experimental accessibility compared to the cleaner but statistically limited dilepton final state. However, this channel also faces challenges, particularly in the identification of jet flavour and the assignment of jets to partons. Nevertheless, the strategies developed in these analyses, ranging from methods used to identify spin analysers to the use of machine learning techniques for flavour-tagging and $t\bar{t}$ reconstruction, provide a strong foundation for future measurements in this area.

Looking ahead, the ongoing Run-3 data-taking period is expected to deliver a substantial increase in integrated luminosity, opening the door to more precise measurements with improved sensitivity. In particular, the analysis in the boosted regime would significantly benefit from the inclusion of

Run-3 data, leading to improved statistical precision. Advances in flavour-tagging algorithms, such as the graph-neural network based GN2 algorithm used in this thesis, are expected to enhance the identification of down-type jets, while further developments in machine learning-based reconstruction methods, building on frameworks such as SPANet, can improve the reconstruction of the $t\bar{t}$ system and jet assignment. Ensuring consistency between particle-level and reconstruction-level definitions in the threshold entanglement measurement, for instance by harmonising reconstruction strategies, would also help reduce potential biases and improve sensitivity.

In the boosted regime, the high-luminosity upgrade of the LHC is expected to play a transformative role by significantly mitigating the limitations arising from low statistics at very high $m_{t\bar{t}}$. To maintain entanglement sensitivity in this extreme kinematic regime, an alternative reconstruction strategy can be adopted in which the hadronically decaying top quark is reconstructed as a single large-R jet, rather than using RC jets reconstructed from fixed-size inputs. This approach naturally captures the highly collimated decay products of the boosted top quark within a single jet and provides access to jet substructure observables that are directly sensitive to the top quark decay. Developing dedicated calibrations for the jet and substructure observables of such large-R jets remains an important area for further study. Along with this, the application of machine-learning-based taggers targeting specific decay patterns (such as charm-enriched decays) could significantly improve the performance of this measurement.

Alongside developments in software and analysis techniques, a crucial aspect that warrants further investigation is the theoretical modelling itself. In particular, a better understanding of parton shower uncertainties remains an important challenge, especially in resolving discrepancies between different generator predictions that contribute significantly in the $t\bar{t}$ threshold region. For example, ongoing efforts to replace the Pythia versus Herwig comparison with variations of parameters within a single generator could provide a more robust definition of this uncertainty. Beyond this, continued efforts to improve the modelling of the threshold regime are essential. In the context of $t\bar{t}$ quasi-bound-state effects, this includes extending the MC models to incorporate additional NRQCD effects, beyond S-wave contributions, including higher J-states and octet contributions. Improved matching between the NRQCD and pQCD contributions of $t\bar{t}$ will also be beneficial in accurately describing the effects at threshold [260]. Furthermore, exploring models that allow for potential BSM effects, or interference between quasi-bound states and SM+BSM $t\bar{t}$ production, represents an important direction for future studies [261]. More broadly, ongoing efforts to refine pQCD predictions for $t\bar{t}$ production are expected to play a significant role. One such example is the bb4l model [245], which has already been implemented in threshold measurements in the dilepton channel [21]. This model includes off-shell and non-resonant contributions, as well as exact spin correlations at NLO, and is expected to simulate the full $b\bar{b}q\bar{q}'\nu\bar{\nu}$ final state in the context of the lepton+jets channel. Another promising approach is based on the MiNNLO framework, which provides NNLO-accurate event generation matched to parton showers with all-order radiative corrections [262]. Such predictions have the potential to enhance the description of $t\bar{t}$ production near threshold, thereby improving our understanding of the underlying dynamics in this region.

The exploration of quantum mechanical properties in the top-quark sector does not end with entanglement. The mixed quantum state of the $t\bar{t}$ system gives access to a broader hierarchy of quantum correlations [263], as illustrated in Figure 9.1. At the most general level, the presence of spin correlations already encodes information about the underlying spin dynamics for separable $t\bar{t}$ states. An alternative definition of quantumness is quantum discord, which is a basic form of quantum correlations and can exhibit asymmetry between subsystems, making it a potentially sensitive probe of

effects such as CP violation [264]. While all entangled states possess non-zero discord, the converse is not necessarily true. Progressing further along this hierarchy, beyond entanglement, one encounters steering. This characterises the ability of one subsystem to non-locally affect or "steer" the state of the other through measurements. At the strongest level of correlations lie the Bell correlated states which are expected to violate Bell inequalities if a loophole-free Bell test were possible in a collider environment [86–88]. Some of these measurements are challenging to perform in conventional low-energy laboratory systems, but are naturally accessible at the LHC. The quantum observables developed for these studies also open the door to probing additional quantum resources, such as "magic," which quantifies the non-classicality required for quantum computational advantage [265]. Evidence for such effects has already been reported by CMS [266]. Alongside these correlations, it is equally important to consider processes that degrade them. Quantum decoherence, arising from interactions with the environment, can reduce or destroy correlations. In the $t\bar{t}$ system, such effects may be induced, for example, by final-state radiation. Recent studies [267] highlight the potential to probe these effects, offering a more complete picture of quantum dynamics in high-energy collisions.

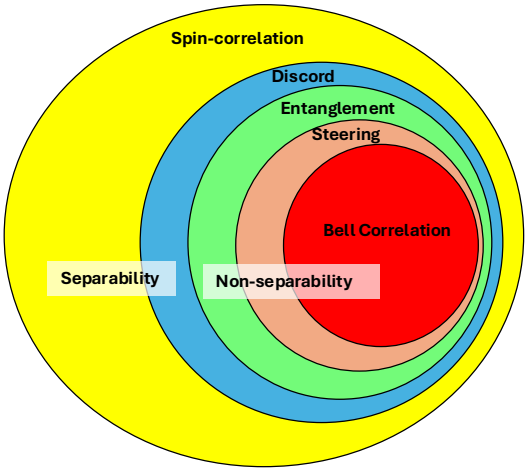


Figure 9.1: Schematic illustration of the relation between the different quantum correlations.

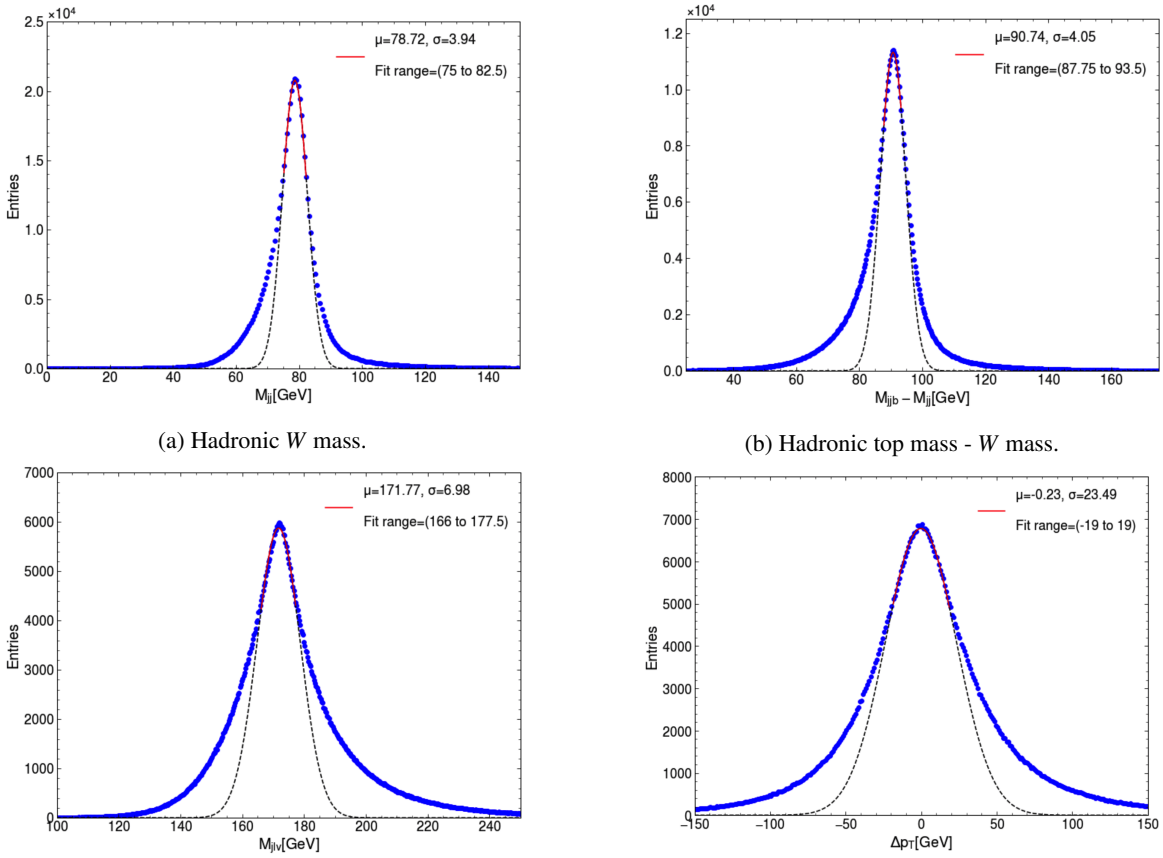
Beyond the top quark sector, quantum information concepts are beginning to be explored in other systems. A recent example is the study of $h \to ZZ^*$ decays by CMS [268], where the spin-1 nature of Z bosons provides a system of qutrits rather than qubits. Additionally, theoretical work [269] suggests that similar quantum effects could be probed in bottom-quark pair production. These developments point to a rapidly growing interface between quantum information science and high-energy physics.

Given that many of these quantum phenomena have traditionally been studied outside the high-energy physics domain, it is essential to clearly communicate the unique opportunities offered by collider experiments. Contributions to initiatives such as the European Strategy for Particle Physics [270] emphasise the potential of colliders to advance quantum information studies in previously unexplored regimes. Together, these efforts promise not only to deepen our understanding of quantum theory in the high-energy frontier but also to provide new avenues for probing the SM and searching for physics beyond it.

APPENDIX A

Kinematic fits for the inputs to the particle-level χ^2 reconstruction

This appendix presents the Gaussian fits for the parameters entering the χ^2 function defined in Equation 7.2 of Section 7.5, as shown in Figures A.1(a)-(d). The parameters are extracted using the from the pQCD $t\bar{t}$ samples at particle-level, using events which have jets that are matched to parton-level objects within a cone of $\Delta R < 0.3$, and satisfying $m_{t\bar{t}} < 600\,\text{GeV}$. The fit regions (shown in red) are defined to capture the central, Gaussian-like portion of each distribution, thereby ensuring accurate modelling of the peak region.



(a) Hadronic W mass.

(b) Hadronic top mass - W mass.

(c) Leptonic top mass.

(d) Scalar difference of the p_{T} of the top quarks.

Figure A.1: Fits of various kinematic parameters entering the χ^2 algorithm.

APPENDIX B

Optimisation studies for entanglement signal region

This appendix outlines the studies conducted to determine the optimal selections on $m_{t\bar{t}}$ and $\beta_{t\bar{t}}$ for defining the low-$m_{t\bar{t}}$ entanglement signal region described in Section 7.6. The signal used in these studies includes pQCD $t\bar{t}$ and the simplified η_t model to account for $t\bar{t}$ quasi-bound-state effects. Various scenarios are tested, and for each, the value of the entanglement observable D_{particle} and the corresponding entanglement limit for POWHEG+PYTHIA, $D^{-1/3}_{\text{particle}}$, are listed in Table B.1. It also includes the expected signal significance, $Z(Pythia + \eta_t)$, calculated using Equation 7.6 for each selection. It should be noted that these studies are performed using a subset of systematic uncertainties, specifically those related to PS and hadronisation, PS/ME matching, recoil-to-top, underlying event, colour reconnection, and top-quark mass.

Table B.1: Summary of results obtained using different selection cuts to optimise the entanglement signal region.

Selection	$D_{\text{particle}} \pm \text{stat} \pm \text{syst}$	$D^{-1/3}_{\text{particle}} \pm \text{syst}$	$Z(Pythia + \eta_t)$
$m_{t\bar{t}} < 390\,\text{GeV}, \beta_{t\bar{t}} < 0.55$	$-0.362 \pm 0.008^{+0.022}_{-0.029}$	-0.228 ± 0.006	5.68
$m_{t\bar{t}} < 390\,\text{GeV}, \beta_{t\bar{t}} < 0.60$	$-0.361 \pm 0.007^{+0.027}_{-0.029}$	-0.229 ± 0.006	4.75
$m_{t\bar{t}} < 390\,\text{GeV}, \beta_{t\bar{t}} < 0.50$	$-0.365 \pm 0.007^{+0.032}_{-0.036}$	-0.228 ± 0.006	4.13
$m_{t\bar{t}} < 390\,\text{GeV}$	$-0.340 \pm 0.005 \pm 0.027$	-0.244 ± 0.004	3.54
$m_{t\bar{t}} < 380\,\text{GeV}, \beta_{t\bar{t}} < 0.55$	$-0.380 \pm 0.009^{+0.027}_{-0.025}$	-0.228 ± 0.004	5.43
$m_{t\bar{t}} < 380\,\text{GeV}, \beta_{t\bar{t}} < 0.60$	$-0.377 \pm 0.008^{+0.030}_{-0.029}$	-0.228 ± 0.003	4.74
$m_{t\bar{t}} < 380\,\text{GeV}, \beta_{t\bar{t}} < 0.50$	$-0.382 \pm 0.009 \pm 0.033$	-0.228 ± 0.004	4.52
$340 < m_{t\bar{t}} < 380\,\text{GeV}, \beta_{t\bar{t}} < 0.55$	$-0.373 \pm 0.010^{+0.038}_{-0.034}$	-0.223 ± 0.007	3.83
$340 < m_{t\bar{t}} < 380\,\text{GeV}, \beta_{t\bar{t}} < 0.60$	$-0.370 \pm 0.010^{+0.039}_{-0.036}$	-0.223 ± 0.006	3.61
$340 < m_{t\bar{t}} < 380\,\text{GeV}, \beta_{t\bar{t}} < 0.50$	$-0.375 \pm 0.011^{+0.041}_{-0.040}$	-0.223 ± 0.007	3.57
$340 < m_{t\bar{t}} < 380\,\text{GeV}$	$-0.351 \pm 0.007^{+0.047}_{-0.044}$	-0.240 ± 0.007	2.32

As shown in the table, applying a cut on $\beta_{t\bar{t}}$ generally improves the expected signal significance. In contrast, imposing a lower bound on $m_{t\bar{t}}$ (similar to the one used in the dilepton analysis), tends to reduce the significance by a noticeable margin. Based on the highest expected significance, the selection $m_{t\bar{t}} < 390\,\text{GeV}, \beta_{t\bar{t}} < 0.55$ is selected as the optimal choice.

APPENDIX C

Unfolding inputs in muon region

In this appendix, the migration, acceptance, efficiency, and response matrix histograms are shown for the SR1 and SR2 corersponding to $\mu + jets$. They exhibit similar behaviour as those for e+jets shown in Figures 8.3 and 8.4.

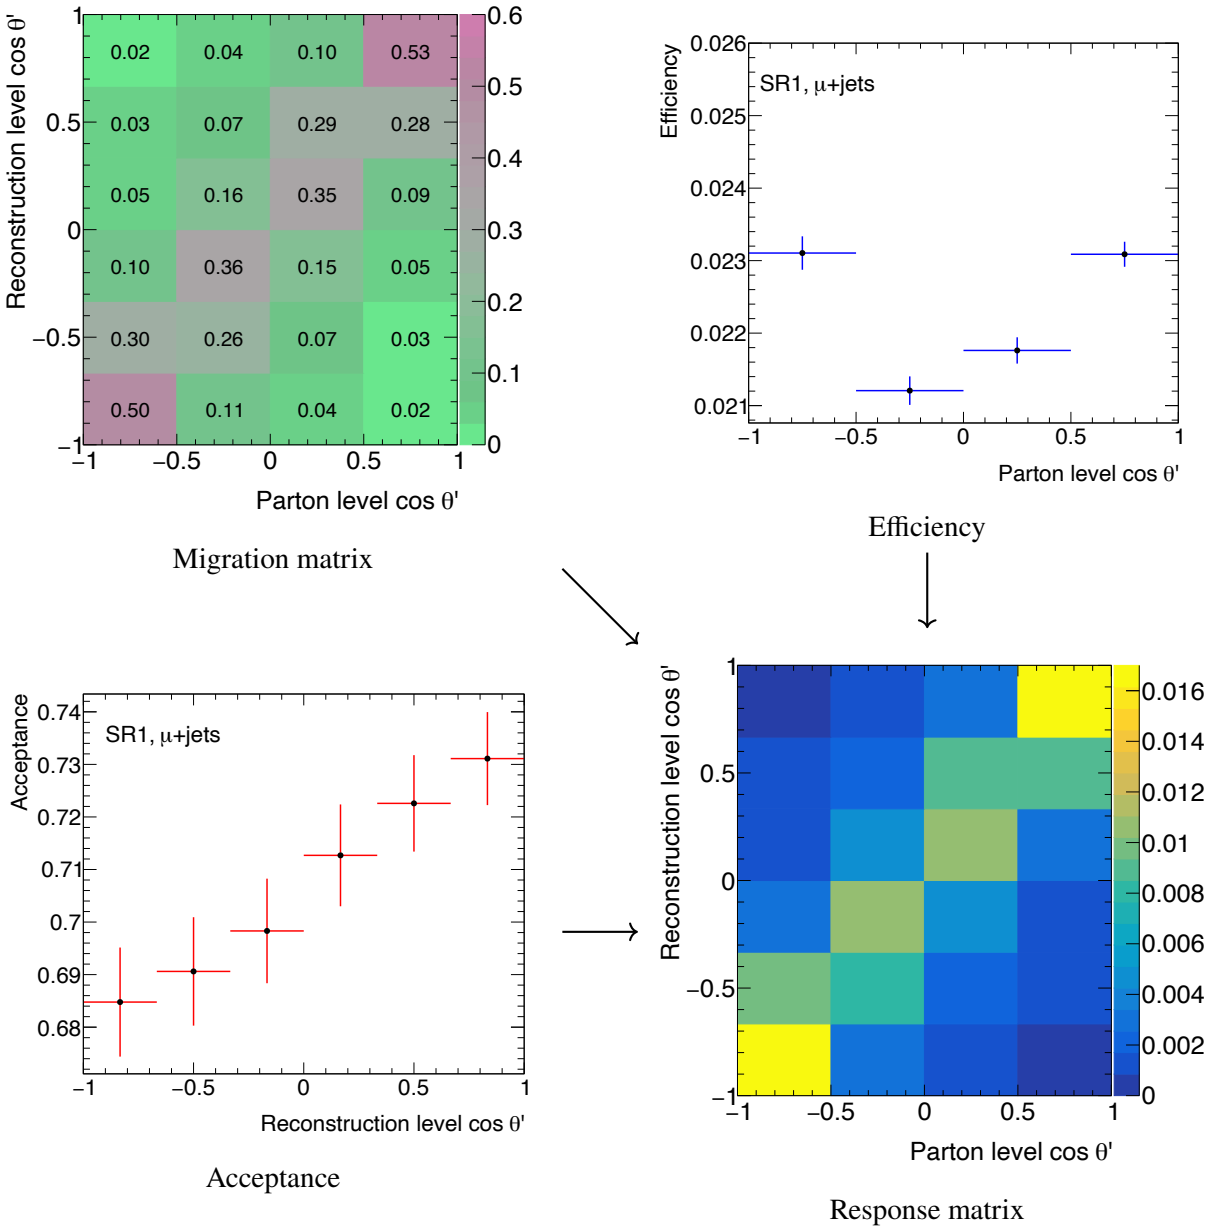


Figure C.1: Migration matrix, acceptance, efficiency, and the resulting response matrix SR1 μ+jets. The z-axis of the migration matrix denotes the probability that an event reconstructed in a given reconstruction-level bin originates from a particular parton-level bin. These values are explicitly shown in the matrix itself. For the response matrix, the corresponding axis represents the same probability after accounting for detector acceptance and efficiency effects.

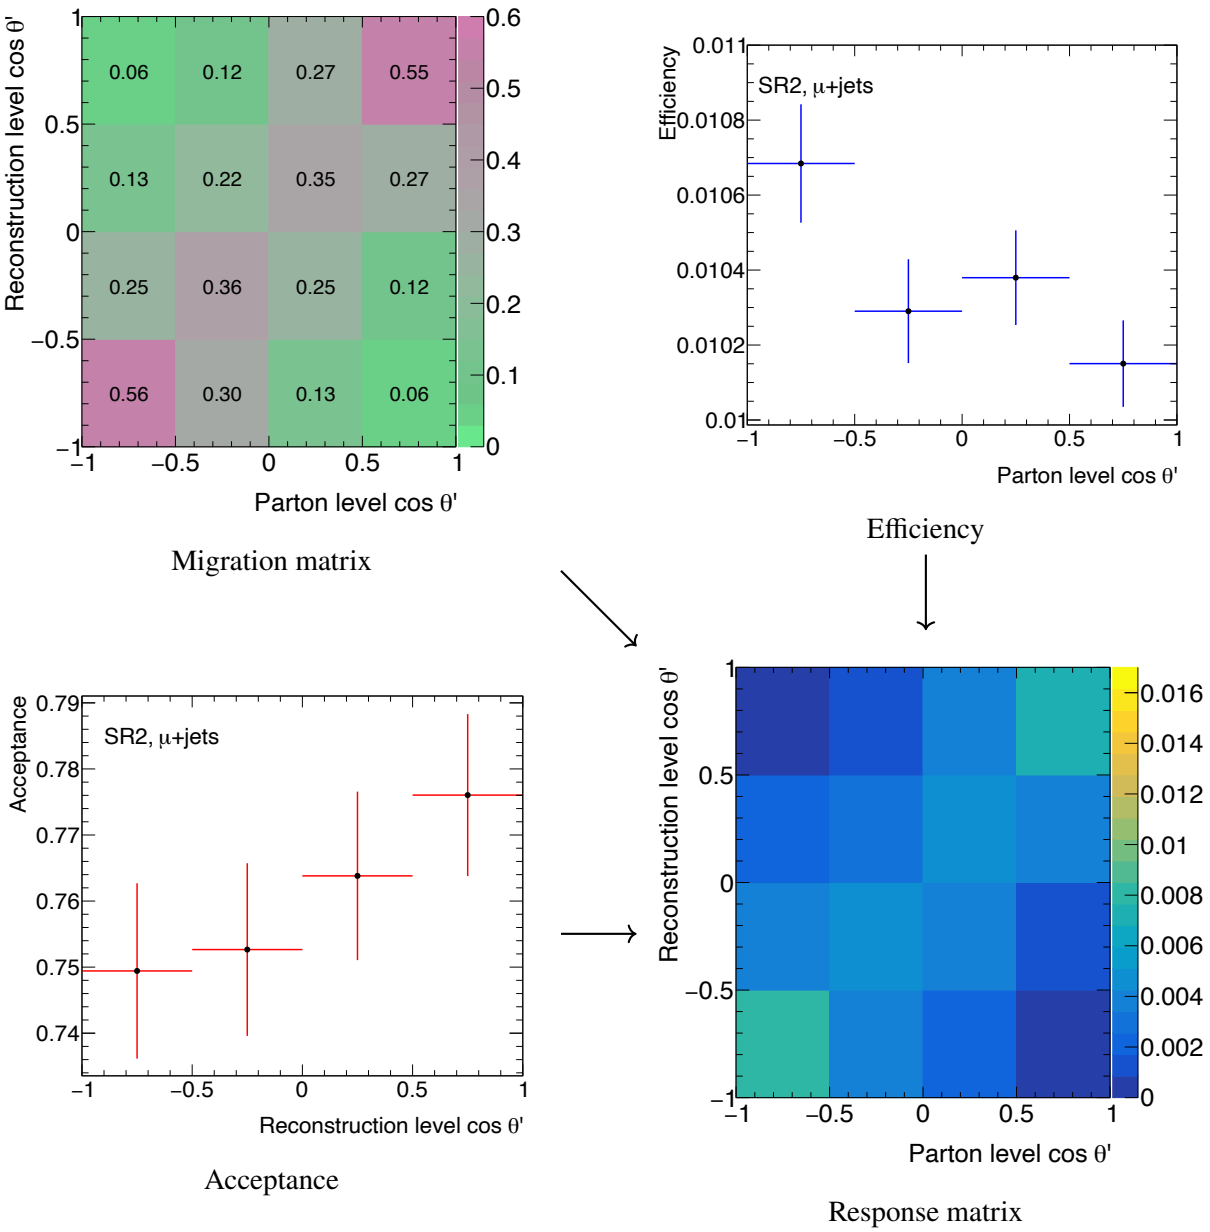


Figure C.2: Migration matrix, acceptance, efficiency, and the resulting response matrix SR2 μ+jets. The z-axis of the migration matrix denotes the probability that an event reconstructed in a given reconstruction-level bin originates from a particular parton-level bin. These values are explicitly shown in the matrix itself. For the response matrix, the corresponding axis represents the same probability after accounting for detector acceptance and efficiency effects.

Bibliography

[1] P. S. E. R. F. R. S.,
LIV. Collision of α particles with light atoms. IV. An anomalous effect in nitrogen,
Phil. Mag. Ser. 6 **37** (1919) 581 (cit. on p. 2).

[2] J. Chadwick, *The Existence of a Neutron*, Proc. Roy. Soc. Lond. A **136** (1932) 692
(cit. on p. 2).

[3] E. D. Bloom et al., *High-Energy Inelastic $e - p$ Scattering at 6° and 10°*,
Phys. Rev. Lett. **23** (16 1969) 930 (cit. on p. 2).

[4] M. Breidenbach et al., *Observed Behavior of Highly Inelastic Electron-Proton Scattering*,
Phys. Rev. Lett. **23** (16 1969) 935 (cit. on p. 2).

[5] TASSO Collaboration, *Evidence for planar events in e+e- annihilation at high energies*,
Physics Letters B **86** (1979) 243, ISSN: 0370-2693 (cit. on p. 2).

[6] M. Banner et al., *Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN pp collider*,
Physics Letters B **122** (1983) 476, ISSN: 0370-2693 (cit. on p. 2).

[7] G. Arnison et al., *Experimental observation of lepton pairs of invariant mass around 95 GeV/c2 at the CERN SPS collider*, Physics Letters B **126** (1983) 398, ISSN: 0370-2693
(cit. on p. 2).

[8] D0 Collaboration, *Observation of the top quark*, Phys. Rev. Lett. **74** (1995) 2632,
arXiv: hep-ex/9503003 (cit. on pp. 2, 40, 45).

[9] CDF Collaboration,
Observation of Top Quark Production in $\overline{p}p$ Collisions with the Collider Detector at Fermilab,
Phys. Rev. Lett. **74** (14 1995) 2626 (cit. on pp. 2, 40, 45).

[10] ATLAS Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Physics Letters B **716** (2012) 1,
ISSN: 0370-2693 (cit. on pp. 2, 17).

[11] CMS Collaboration,
Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC,
Physics Letters B **716** (2012) 30, ISSN: 0370-2693 (cit. on pp. 2, 17).

[12] J. S. Bell, *On the Einstein Podolsky Rosen paradox*, Physics Physique Fizika **1** (3 1964) 195
(cit. on pp. 3, 21, 38).

[13] A. Einstein, B. Podolsky and N. Rosen,
Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?,
Phys. Rev. **47** (10 1935) 777 (cit. on pp. 3, 21, 38).

[14] A. Aspect, P. Grangier and G. Roger, *Experimental Realization of Einstein-Podolsky-Rosen-Bohm Gedankenexperiment: A New Violation of Bell's Inequalities*,
Phys. Rev. Lett. **49** (2 1982) 91 (cit. on pp. 3, 21).

[15] E. Hagley et al., *Generation of Einstein-Podolsky-Rosen Pairs of Atoms*,
Phys. Rev. Lett. **79** (1 1997) 1 (cit. on pp. 3, 21).

[16] M. Steffen et al.,
Measurement of the Entanglement of Two Superconducting Qubits via State Tomography,
Science **313** (2006) 1423 (cit. on pp. 3, 21).

[17] ATLAS Collaboration,
Observation of quantum entanglement with top quarks at the ATLAS detector,
Nature **633** (2024), ISSN: 1476-4687 (cit. on pp. 3, 37, 40, 97, 115, 131, 133, 139, 179).

[18] V. S. Fadin and V. A. Khoze,
Threshold Behavior of Heavy Top Production in e+ e- Collisions, JETP Lett. (1987)
(cit. on pp. 3, 40).

[19] V. S. Fadin, V. A. Khoze and T. Sjostrand,
On the Threshold Behavior of Heavy Top Production, Z. Phys. C **48** (1990) 613
(cit. on pp. 3, 40, 70).

[20] Y. Sumino, K. Fujii, K. Hagiwara, H. Murayama and C.-K. Ng,
Top-quark pair production near threshold, Phys. Rev. D (1993) (cit. on pp. 3, 40, 42).

[21] ATLAS Collaboration, *Observation of a cross-section enhancement near the $t\bar{t}$ production threshold in $\sqrt{s} = 13\,TeV$ pp collisions with the ATLAS detector*,
Reports on Progress in Physics **89** (2026) 057801,
URL: https://doi.org/10.1088/1361-6633/ae60a0
(cit. on pp. 4, 40, 43, 67, 125, 136, 180).

[22] CMS Collaboration,
Observation of a pseudoscalar excess at the top quark pair production threshold,
Reports on Progress in Physics **88** (2025) 087801,
URL: https://doi.org/10.1088/1361-6633/adf7d3 (cit. on pp. 4, 40, 43, 125, 136).

[23] CMS Collaboration, *Observation of quantum entanglement in top quark pair production in proton–proton collisions at $\sqrt{s} = 13$ TeV*, Reports on Progress in Physics **87** (2024)
(cit. on pp. 4, 37, 139, 140, 179).

[24] CMS Collaboration,
Measurements of polarization and spin correlation and observation of entanglement in top quark pairs using lepton + *jets events from proton–proton collisions at $\sqrt{s} = 13$ TeV*,
Phys. Rev. D **110** (11 2024) (cit. on pp. 4, 34, 37, 140, 141, 176, 178, 179).

[25] ATLAS Collaboration, *Search for ttbar resonances in final states with exactly one or two leptons using 140 fb^{-1} of pp collision data at $\sqrt{s} = 13$ TeV with the ATLAS experiment*,
(2025), arXiv: 2512.17856 [hep-ex] (cit. on pp. 5, 107).

[26] ATLAS Collaboration, *Track and Vertex Reconstruction with the ATLAS Inner Detector*, (2026), arXiv: 2605.07585 [physics.ins-det] (cit. on pp. 5, 72, 73, 75–78).

[27] F. Halzen and A. D. Martin, *Quarks and Leptons: An Introductory Course in Modern Particle Physics*, Wiley, 1984, ISBN: 9780471887416 (cit. on pp. 7, 23).

[28] M. Thomson, *Modern particle physics*, Cambridge University Press, 2013 (cit. on p. 7).

[29] M. E. Peskin and D. V. Schroeder, *An Introduction to quantum field theory*, Reading, USA: Addison-Wesley, 1995 (cit. on p. 7).

[30] O. Stern, *Eine direkte Messung der thermischen Molekulargeschwindigkeit*, Zeitschrift für Physik **2** (1920) 49, URL: https://api.semanticscholar.org/CorpusID:123548450 (cit. on p. 8).

[31] Particle Data Group, *Review of particle physics*, Phys. Rev. D **110** (2024) 030001 (cit. on pp. 8, 9, 18, 21, 22, 26, 82, 90, 105, 111).

[32] Wikipedia contributors, *Standard Model — Wikipedia, The Free Encyclopedia*, https://en.wikipedia.org/w/index.php?title=Standard_Model&oldid=1108241997, [Online; accessed 09-November-2025], 2025 (cit. on p. 9).

[33] H. Georgi, *Lie algebras in particle physics: from isospin to unified theories*, Taylor & Francis, 2000 (cit. on p. 10).

[34] P. A. M. Dirac, *The quantum theory of the electron*, Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character **117** (1928) 610, URL: https://royalsocietypublishing.org/doi/abs/10.1098/rspa.1928.0023 (cit. on p. 10).

[35] J. Altmann, A. Dubla, V. Greco, A. Rossi and P. Skands, *Towards the understanding of heavy quarks hadronization: from leptonic to heavy-ion collisions*, The European Physical Journal C **85** (2025), ISSN: 1434-6052 (cit. on p. 12).

[36] S. Weinberg, *A Model of Leptons*, Phys. Rev. Lett. **19** (21 1967) 1264 (cit. on p. 13).

[37] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes and R. P. Hudson, *Experimental Test of Parity Conservation in Beta Decay*, Phys. Rev. **105** (4 1957) 1413 (cit. on p. 13).

[38] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, Phys. Rev. Lett. **13** (9 1964) 321 (cit. on p. 15).

[39] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, Phys. Rev. Lett. **13** (16 1964) 508 (cit. on p. 15).

[40] G. S. Guralnik, C. R. Hagen and T. W. B. Kibble, *Global Conservation Laws and Massless Particles*, Phys. Rev. Lett. **13** (20 1964) 585 (cit. on p. 15).

[41] J. Ellis, M. K. Gaillard and D. V. Nanopoulos, "A Historical Profile of the Higgs Boson", *The Standard Theory of Particle Physics*, chap. Chapter 14 255 (cit. on pp. 16, 17).

[42] N. Cabibbo, *Unitary Symmetry and Leptonic Decays*, Phys. Rev. Lett. **10** (12 1963) 531 (cit. on p. 18).

[43] M. Kobayashi and T. Maskawa, *CP-Violation in the Renormalizable Theory of Weak Interaction*, Progress of Theoretical Physics **49** (1973) 652, ISSN: 0033-068X (cit. on p. 18).

[44] J. F. Donoghue, *The effective field theory treatment of quantum gravity*, AIP Conference Proceedings **1483** (2012) 73, ISSN: 0094-243X (cit. on p. 19).

[45] A. D. Sakharov, *Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe*, Soviet Physics Uspekhi **34** (1991) 392 (cit. on p. 19).

[46] J. H. Christenson, J. W. Cronin, V. L. Fitch and R. Turlay, *Evidence for the 2π Decay of the K_2^0 Meson*, Phys. Rev. Lett. **13** (4 1964) 138 (cit. on p. 19).

[47] The T2K Collaboration, *Constraint on the matter-antimatter symmetry-violating phase in neutrino oscillations*, Nature **580** (2020), ISSN: 1476-4687 (cit. on p. 19).

[48] V. C. Rubin and J. Ford W. Kent, *Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions*, The Astrophysical Journal **159** (1970) 379 (cit. on p. 19).

[49] Planck Collaboration et al., *Planck 2015 results - XIII. Cosmological parameters*, Astronomy & Astrophysics **594** (2016) A13 (cit. on p. 19).

[50] R. Massey, T. Kitching and J. Richard, *The dark matter of gravitational lensing*, Reports on Progress in Physics **73** (2010) 086901 (cit. on p. 19).

[51] Planck Collaboration, *Planck 2013 results. I. Overview of products and scientific results*, Astron. Astrophys. **571** (2014) A1, arXiv: 1303.5062 [astro-ph.CO] (cit. on p. 19).

[52] P. D. Bolton, F. F. Deppisch and P. S. B. Dev, *Probes of Heavy Sterile Neutrinos*, 2022, arXiv: 2206.01140 [hep-ph] (cit. on p. 19).

[53] C. B. Adams et al., *Axion Dark Matter*, 2023, arXiv: 2203.14923 [hep-ex] (cit. on p. 19).

[54] L. Roszkowski, E. M. Sessolo and S. Trojanowski, *WIMP dark matter candidates and searches–current status and future prospects*, Reports on Progress in Physics **81** (2018) 066201 (cit. on p. 19).

[55] Z. Maki, M. Nakagawa and S. Sakata, *Remarks on the Unified Model of Elementary Particles*, Progress of Theoretical Physics **28** (1962) 870, ISSN: 0033-068X (cit. on p. 19).

[56] Super-Kamiokande Collaboration, *Evidence for Oscillation of Atmospheric Neutrinos*, Phys. Rev. Lett. **81** (8 1998) 1562 (cit. on p. 19).

[57] SNO Collaboration, *Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by ^8B Solar Neutrinos at the Sudbury Neutrino Observatory*, Phys. Rev. Lett. **87** (7 2001) 071301 (cit. on p. 19).

[58] E. Majorana, *Teoria simmetrica dell'elettrone e del positrone*, Nuovo Cim. **14** (1937) 171 (cit. on p. 19).

[59] G. C. Branco et al., *Theory and phenomenology of two-Higgs-doublet models*, Phys. Rept. **516** (2012) 1, arXiv: 1106.0034 [hep-ph] (cit. on p. 20).

[60] H. Haber and G. Kane, *The search for supersymmetry: Probing physics beyond the standard model*, Physics Reports **117** (1985) 75, ISSN: 0370-1573, URL: https://www.sciencedirect.com/science/article/pii/0370157385900511 (cit. on p. 20).

[61] D0 Collaboration, *Improved determination of the width of the top quark*, Phys. Rev. D **85** (9 2012) 091104 (cit. on p. 21).

[62] I. Bigi, Y. Dokshitzer, V. Khoze, J. Kühn and P. Zerwas, *Production and decay properties of ultra-heavy quarks*, Physics Letters B **181** (1986) 157, ISSN: 0370-2693 (cit. on p. 21).

[63] J. C. Collins, D. E. Soper and G. Sterman, *Factorization of Hard Processes in QCD*, 2004, arXiv: hep-ph/0409313 [hep-ph] (cit. on pp. 22, 23).

[64] F. D. Aaron et al., *Combined measurement and QCD analysis of the inclusive $e \pm p$ scattering cross sections at HERA*, Journal of High Energy Physics **2010** (2010), ISSN: 1029-8479 (cit. on p. 23).

[65] J. Butterworth et al., *PDF4LHC recommendations for LHC Run II*, Journal of Physics G: Nuclear and Particle Physics **43** (2016) 023001, ISSN: 1361-6471 (cit. on pp. 23, 24, 122).

[66] J. Gao, L. Harland-Lang and J. Rojo, *The structure of the proton in the LHC precision era*, Physics Reports **742** (2018), ISSN: 0370-1573 (cit. on p. 24).

[67] D0 Collaboration, *Measurement of the Top Quark Mass Using Dilepton Events*, Physical Review Letters **80** (1998), ISSN: 1079-7114 (cit. on p. 25).

[68] R. F. Werner, *Quantum states with Einstein-Podolsky-Rosen correlations admitting a hidden-variable model*, Phys. Rev. A **40** (8 1989) 4277 (cit. on p. 26).

[69] G. Mahlon and S. Parke, *Angular correlations in top quark pair production and decay at hadron colliders*, Physical Review D **53** (1996), ISSN: 1089-4918 (cit. on pp. 27, 28).

[70] A. Brandenburg, Z. Si and P. Uwer, *QCD-corrected spin analysing power of jets in decays of polarized top quarks*, Physics Letters B **539** (2002) 235, ISSN: 0370-2693 (cit. on pp. 28–30, 32, 33).

[71] W. Bernreuther, *Top quark physics at the LHC*, 2008, arXiv: 0805.1333 [hep-ph] (cit. on p. 30).

[72] F. Hubaut, E. Monnier, P. Pralavorio, K. Smolek and V. Simak, *ATLAS sensitivity to top quark and W boson polarization in $t\bar{t}$ events*, Eur. Phys. J. C **44S2** (2005) 13, arXiv: hep-ex/0508061 (cit. on p. 30).

[73] B. Tweedie, *Better hadronic top quark polarimetry*, Phys. Rev. D **90** (9 2014) 094010 (cit. on pp. 30–32).

[74] B. Lemmer, *Spin Correlations in $t\bar{t}$ events from pp Collisions: Measured at $\sqrt{s} = 7$ TeV in the lepton+jets final state with the ATLAS Detector*, Springer, 2015 (cit. on p. 33).

[75] Y. Afik and J. R. M. de Nova,
Entanglement and quantum tomography with top quarks at the LHC,
The European Physical Journal Plus **136** (2021), ISSN: 2190-5444 (cit. on pp. 33–36).

[76] ATLAS Collaboration, *Measurements of top quark spin observables in $t\bar{t}$ events using dilepton final states in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector*,
Journal of High Energy Physics **2017** (2017), ISSN: 1029-8479 (cit. on p. 34).

[77] CMS Collaboration, *Measurement of the top quark polarization and $t\bar{t}$ spin correlations using dilepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV*,
Phys. Rev. D **100** (2019) 072002, ISSN: 2470-0029 (cit. on p. 34).

[78] ATLAS Collaboration, *Measurements of top-quark pair spin correlations in the $e\mu$ channel at $\sqrt{s} = 13$ TeV using pp collisions in the ATLAS detector*,
The European Physical Journal C **80** (2020), ISSN: 1434-6052 (cit. on p. 34).

[79] ATLAS Collaboration, *Combination of ATLAS and CMS measurements of the azimuthal spin correlation variable in dilepton $t\bar{t}$ events at $\sqrt{s} = 13$ TeV*, tech. rep., CERN, 2026,
URL: https://cds.cern.ch/record/2952286 (cit. on p. 34).

[80] ATLAS Collaboration, *Measurements of spin correlation in top-antitop quark events from proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*,
Phys. Rev. D **90** (11 2014) 112016 (cit. on p. 34).

[81] A. Peres, *Separability Criterion for Density Matrices*, Phys. Rev. Lett. **77** (8 1996) (cit. on p. 34).

[82] P. Horodecki,
Separability criterion and inseparable mixed states with positive partial transposition,
Physics Letters A **232** (1997) 333, ISSN: 0375-9601 (cit. on p. 34).

[83] W. K. Wootters, *Entanglement of Formation of an Arbitrary State of Two Qubits*,
Phys. Rev. Lett. **80** (10 1998) 2245 (cit. on p. 35).

[84] T. Han, M. Low and T. A. Wu,
Quantum Entanglement and Bell Inequality Violation in Semi-Leptonic Top Decays, 2023,
arXiv: 2310.17696 [hep-ph] (cit. on pp. 37, 38).

[85] J. F. Clauser, M. A. Horne, A. Shimony and R. A. Holt,
Proposed Experiment to Test Local Hidden-Variable Theories,
Phys. Rev. Lett. **23** (15 1969) 880 (cit. on p. 38).

[86] S. Abel, M. Dittmar and H. Dreiner, *Testing locality at colliders via Bell's inequality?*,
Physics Letters B **280** (1992) 304, ISSN: 0370-2693 (cit. on pp. 39, 181).

[87] P. Bechtle, C. Breuning, H. K. Dreiner and C. Duhr, *A critical appraisal of tests of locality and of entanglement versus non-entanglement at colliders*, 2025,
arXiv: 2507.15947 [hep-ph], URL: https://arxiv.org/abs/2507.15947
(cit. on pp. 39, 181).

[88] S. A. Abel, H. K. Dreiner, R. Sengupta and L. Ubaldi, *Colliders are Testing neither Locality via Bell's Inequality nor Entanglement versus Non-Entanglement*, 2025, arXiv: 2507.15949 [hep-ph], URL: https://arxiv.org/abs/2507.15949 (cit. on pp. 39, 181).

[89] H. N. Russell and F. A. Saunders, *New Regularities in the Spectra of the Alkaline Earths*, The Astrophysical Journal **61** (1925) 38 (cit. on p. 41).

[90] Y. Sumino and H. Yokoya, *Bound-state effects on kinematical distributions of top quarks at hadron colliders*, JHEP **09** (2010) 034, [Erratum: JHEP 06, 037 (2016)], arXiv: 1007.0075 [hep-ph] (cit. on pp. 41, 42, 70).

[91] K. Hagiwara, Y. Sumino and H. Yokoya, *Bound-state effects on top quark production at hadron colliders*, Physics Letters B **666** (2008), ISSN: 0370-2693 (cit. on p. 41).

[92] Y. Kiyo, J. H. Kühn, S. Moch, M. Steinhauser and P. Uwer, *Top-quark pair production near threshold at LHC*, The European Physical Journal C **60** (2009), ISSN: 1434-6052 (cit. on p. 41).

[93] M. V. Garzelli, G. Limatola, S. .-O. Moch, M. Steinhauser and O. Zenaiev, *Updated predictions for toponium production at the LHC*, 2025, arXiv: 2412.16685 [hep-ph], URL: https://arxiv.org/abs/2412.16685 (cit. on p. 41).

[94] F. Maltoni, C. Severi, S. Tentori and E. Vryonidou, *Quantum detection of new physics in top-quark pair production at the LHC*, 2024, arXiv: 2401.08751 [hep-ph], URL: https://arxiv.org/abs/2401.08751 (cit. on pp. 41, 70).

[95] B. Fuks, K. Hagiwara, K. Ma and Y.-J. Zheng, *Simulating toponium formation signals at the LHC*, Eur. Phys. J. C **85** (2025) 157, arXiv: 2411.18962 [hep-ph] (cit. on pp. 43, 69, 70, 179).

[96] L. Evans and P. Bryant, *LHC Machine*, Journal of Instrumentation **3** (2008) S08001, URL: https://doi.org/10.1088/1748-0221/3/08/S08001 (cit. on p. 45).

[97] *CERN images gallery*, https://home.cern/resources/image/cern/views-cern-images-gallery, [Online; accessed 13-November-2025] (cit. on p. 46).

[98] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, Journal of Instrumentation **3** (2008) S08003 (cit. on pp. 47, 50, 53, 56).

[99] CMS Collaboration, *The CMS Experiment at the CERN LHC*, JINST **3** (2008) S08004 (cit. on p. 47).

[100] ALICE Collaboration, *The ALICE experiment at the CERN LHC*, JINST **3** (2008) S08002, Also published by CERN Geneva in 2010, URL: https://cds.cern.ch/record/1129812 (cit. on p. 47).

[101] LHCb Collaboration, *The LHCb Detector at the LHC*, JINST **3** (2008) S08005 (cit. on p. 47).

[102] ATLAS Collaboration, *Luminosity determination in pp collisions at* $\sqrt{s} = 13$ TeV *using the ATLAS detector at the LHC*, The European Physical Journal C **83** (2023), ISSN: 1434-6052 (cit. on pp. 48, 49).

[103] G. Avoni et al.,
The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS,
Journal of Instrumentation **13** (2018) P07017 (cit. on p. 48).

[104] V. Cindro et al., *The ATLAS Beam Conditions Monitor*,
Journal of Instrumentation **3** (2008) P02004 (cit. on p. 48).

[105] S. van der Meer, *Calibration of the effective beam height in the ISR*, tech. rep., CERN, 1968, URL: https://cds.cern.ch/record/296752 (cit. on p. 48).

[106] *Public ATLAS Luminosity Public Results*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResults, Accessed: 2026-05-18, 2026 (cit. on p. 49).

[107] ATLAS Collaboration,
ATLAS data quality operations and performance for 2015–2018 data-taking,
Journal of Instrumentation **15** (2020) P04003,
URL: https://doi.org/10.1088/1748-0221/15/04/P04003 (cit. on p. 49).

[108] ATLAS Collaboration, *The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3*,
Journal of Instrumentation **19** (2024) P05063, ISSN: 1748-0221 (cit. on pp. 50, 53, 54, 57–59).

[109] Bianchi, Riccardo Maria and ATLAS Collaboration,
"ATLAS experiment schematic or layout illustration", General Photo, 2022,
URL: https://cds.cern.ch/record/2837191 (cit. on p. 51).

[110] G. Aad et al., *ATLAS pixel detector electronics and sensors*,
Journal of Instrumentation **3** (2008) P07007 (cit. on p. 52).

[111] ATLAS IBL Collaboration, *Production and integration of the ATLAS Insertable B-Layer*,
Journal of Instrumentation **13** (2018), ISSN: 1748-0221 (cit. on p. 52).

[112] ATLAS Collaboration, *Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: tracker properties*, Journal of Instrumentation **12** (2017) P05002 (cit. on p. 53).

[113] V. Tisserand, *Optimisation du détecteur ATLAS pour la recherche du boson de Higgs se désintegrant en deux photons au LHC*, Theses: Université Paris Sud - Paris XI, 1997,
URL: https://theses.hal.science/tel-00544731 (cit. on p. 55).

[114] ATLAS Collaboration,
ATLAS: Detector and physics performance technical design report. Volume 1, (1999)
(cit. on pp. 55–57).

[115] ATLAS Collaboration, *ATLAS magnet system*, Technical design report. ATLAS,
Geneva: CERN, 1997, URL: https://cds.cern.ch/record/338080 (cit. on p. 56).

[116] G. Di Gregorio, "The Large Hadron Collider and the ATLAS Experiment",
Cham: Springer International Publishing, 2023,
URL: https://doi.org/10.1007/978-3-031-20013-7_2 (cit. on p. 57).

[117] ATLAS collaboration, *Operation of the ATLAS trigger system in Run 2*, Journal of Instrumentation **15** (2020) P10004 (cit. on p. 59).

[118] O. Aberle et al., *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*, CERN Yellow Reports: Monographs, Geneva: CERN, 2020, URL: https://cds.cern.ch/record/2749422 (cit. on pp. 59, 60).

[119] ATLAS Collaboration, *The ATLAS Upgrade for the HL-LHC*, tech. rep., CERN, 2025, URL: https://cds.cern.ch/record/2928798 (cit. on p. 60).

[120] ATLAS Collaboration, *ATLAS Phase-II Upgrade Scoping Document*, tech. rep., CERN, 2015, URL: https://cds.cern.ch/record/2055248 (cit. on p. 60).

[121] ATLAS Collaboration, *Technical Design Report for the ATLAS Inner Tracker Pixel Detector*, tech. rep., CERN, 2017, URL: https://cds.cern.ch/record/2285585 (cit. on p. 60).

[122] ATLAS Collaboration, *Technical Design Report for the ATLAS Inner Tracker Strip Detector*, tech. rep., CERN, 2017, URL: https://cds.cern.ch/record/2257755 (cit. on p. 60).

[123] ATLAS Collaboration, *Technical Design Report: A High-Granularity Timing Detector for the ATLAS Phase-II Upgrade*, tech. rep., CERN, 2020, URL: https://cds.cern.ch/record/2719855 (cit. on p. 60).

[124] A. Buckley et al., *General-purpose event generators for LHC physics*, Physics Reports **504** (2011), ISSN: 0370-1573 (cit. on pp. 61, 62).

[125] J. M. Campbell et al., *Event Generators for High-Energy Physics Experiments*, 2025, arXiv: 2203.11110 [hep-ph], URL: https://arxiv.org/abs/2203.11110 (cit. on p. 61).

[126] S. Alioli, P. Nason, C. Oleari and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, Journal of High Energy Physics **2010** (2010), ISSN: 1029-8479 (cit. on pp. 62, 67–69).

[127] Sherpa Collaboration, *Event Generation with Sherpa 2.2*, SciPost Phys. **7** (2019) 034, arXiv: 1905.09127 [hep-ph] (cit. on pp. 62, 69).

[128] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, Journal of High Energy Physics **2014** (2014), ISSN: 1029-8479 (cit. on pp. 62, 69).

[129] R. Frederix et al., *The automation of next-to-leading order electroweak calculations*, Journal of High Energy Physics **2018** (2018), ISSN: 1029-8479 (cit. on p. 62).

[130] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, Computer Physics Communications **191** (2015) 159, ISSN: 0010-4655 (cit. on pp. 62–64, 67).

[131] M. Bähr et al., *Herwig++ physics and manual*, The European Physical Journal C **58** (2008), ISSN: 1434-6052 (cit. on pp. 62, 63, 122).

[132] J. Bellm et al., *Herwig 7.0/Herwig++ 3.0 release note*, The European Physical Journal C **76** (2016), ISSN: 1434-6052 (cit. on pp. 62, 63, 122).

[133] J. Jamieson,
Aiming For The Top Corner: Measuring top-quark pair production in association with additional jets in highly boosted events using the ATLAS detector,CERN-THESIS-2022-025, 2022, URL: https://cds.cern.ch/record/2805713 (cit. on p. 63).

[134] T. Sjöstrand, *Jet fragmentation of multiparton configurations in a string framework*, Nuclear Physics B **248** (1984) 469, ISSN: 0550-3213 (cit. on p. 63).

[135] S. Agostinelli et al., *Geant4—a simulation toolkit*,
Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **506** (2003) 250, ISSN: 0168-9002 (cit. on p. 64).

[136] ATLAS Collaboration, *The ATLAS Simulation Infrastructure*,
The European Physical Journal C **70** (2010) 823, ISSN: 1434-6052,
URL: http://dx.doi.org/10.1140/epjc/s10052-010-1429-9 (cit. on pp. 64, 158).

[137] ATLAS Collaboration, *The Pythia 8 A3 tune description of ATLAS minimum bias and inelastic measurements incorporating the Donnachie-Landshoff diffractive model*, tech. rep., CERN, 2016, URL: https://cds.cern.ch/record/2206965 (cit. on p. 64).

[138] D. J. Lange, *The EvtGen particle decay simulation package*,
Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **462** (2001) 152, BEAUTY2000, Proceedings of the 7th Int. Conf. on B-Physics at Hadron Machines, ISSN: 0168-9002 (cit. on p. 66).

[139] ATLAS Collaboration, *ATLAS Pythia 8 tunes to 7 TeV data*, tech. rep., CERN, 2014,
URL: https://cds.cern.ch/record/1966419 (cit. on pp. 67, 123, 124).

[140] R. D. Ball et al., *Parton distributions with LHC data*, Nuclear Physics B **867** (2013), ISSN: 0550-3213 (cit. on p. 67).

[141] S. Frixione, G. Ridolfi and P. Nason,
A positive-weight next-to-leading-order Monte Carlo for heavy flavour hadroproduction, Journal of High Energy Physics **2007** (2007), ISSN: 1029-8479 (cit. on pp. 67, 69).

[142] P. Nason, *A New Method for Combining NLO QCD with Shower Monte Carlo Algorithms*, Journal of High Energy Physics **2004** (2004), ISSN: 1029-8479 (cit. on pp. 67–69).

[143] S. Frixione, P. Nason and C. Oleari,
Matching NLO QCD computations with parton shower simulations: the POWHEG method, Journal of High Energy Physics **2007** (2007), ISSN: 1029-8479 (cit. on pp. 67–69).

[144] R. D. Ball et al., *Parton distributions for the LHC run II*,
Journal of High Energy Physics **2015** (2015), ISSN: 1029-8479 (cit. on pp. 67–69).

[145] ATLAS Collaboration, *Studies on top-quark Monte Carlo modelling for Top2016*, tech. rep., CERN, 2016, URL: https://cds.cern.ch/record/2216168 (cit. on pp. 67, 68).

[146] M. Beneke, P. Falgari, S. Klein and C. Schwinn,
Hadronic top-quark pair production with NNLL threshold resummation,
Nuclear Physics B **855** (2012), ISSN: 0550-3213 (cit. on p. 67).

[147] M. Cacciari, M. Czakon, M. Mangano, A. Mitov and P. Nason, *Top-pair production at hadron colliders with next-to-next-to-leading logarithmic soft-gluon resummation*, Physics Letters B **710** (2012), ISSN: 0370-2693 (cit. on p. 67).

[148] P. Bärnreuther, M. Czakon and A. Mitov, *Percent-Level-Precision Physics at the Tevatron: Next-to-Next-to-Leading Order QCD Corrections to $q\bar{q} \to t\bar{t} + X$*, Physical Review Letters **109** (2012), ISSN: 1079-7114 (cit. on p. 67).

[149] M. Czakon and A. Mitov, *NNLO corrections to top-pair production at hadron colliders: the all-fermionic scattering channels*, Journal of High Energy Physics **2012** (2012), ISSN: 1029-8479 (cit. on p. 67).

[150] M. Czakon, P. Fiedler and A. Mitov, *Total Top-Quark Pair-Production Cross Section at Hadron Colliders Through $O(\alpha_s^4$*, Physical Review Letters **110** (2013), ISSN: 1079-7114 (cit. on p. 67).

[151] M. Czakon and A. Mitov, *NNLO corrections to top pair production at hadron colliders: the quark-gluon reaction*, Journal of High Energy Physics **2013** (2013), ISSN: 1029-8479 (cit. on p. 67).

[152] M. Czakon and A. Mitov, *Top++: A program for the calculation of the top-pair cross-section at hadron colliders*, Computer Physics Communications **185** (2014), ISSN: 0010-4655 (cit. on p. 67).

[153] M. Grazzini, S. Kallweit and M. Wiesemann, *Fully differential NNLO computations with MATRIX*, The European Physical Journal C **78** (2018), ISSN: 1434-6052 (cit. on p. 67).

[154] M. Aliev et al., *HATHOR – HAdronic Top and Heavy quarks crOss section calculatoR*, Computer Physics Communications **182** (2011), ISSN: 0010-4655 (cit. on p. 67).

[155] J. H. Kühn, A. Scharf and P. Uwer, *Weak interactions in top-quark pair production at hadron colliders: An update*, Physical Review D **91** (2015), ISSN: 1550-2368 (cit. on p. 67).

[156] J. H. Kühn, A. Scharf and P. Uwer, *Electroweak corrections to top-quark pair productionin quark-antiquark annihilation*, The European Physical Journal C **45** (2006), ISSN: 1434-6052 (cit. on p. 67).

[157] S. Frixione, E. Laenen, P. Motylinski, C. White and B. R. Webber, *Single-top hadroproduction in association with aWboson*, Journal of High Energy Physics **2008** (2008), ISSN: 1029-8479 (cit. on p. 68).

[158] C. D. White, S. Frixione, E. Laenen and F. Maltoni, *Isolating Wt production at the LHC*, Journal of High Energy Physics **2009** (2009) 074, URL: https://doi.org/10.1088/1126-6708/2009/11/074 (cit. on p. 68).

[159] S. Frixione, E. Laenen, P. Motylinski and B. R. Webber, *Angular correlations of lepton pairs from vector boson and top quark decays in Monte Carlo simulations*, Journal of High Energy Physics **2007** (2007), ISSN: 1029-8479 (cit. on p. 68).

[160] P. Artoisenet, R. Frederix, O. Mattelaer and R. Rietkerk, *Automatic spin-entangled decays of heavy resonances in Monte Carlo simulations*, Journal of High Energy Physics **2013** (2013), ISSN: 1029-8479 (cit. on p. 68).

[161] N. Kidonakis and N. Yamanaka,
Higher-order corrections for tW production at high-energy hadron colliders,
Journal of High Energy Physics **2021** (2021), ISSN: 1029-8479 (cit. on p. 68).

[162] Z. L. Liu and J. Gao,
s-channel single top quark production and decay at next-to-next-to-leading-order in QCD,
Physical Review D **98** (2018), ISSN: 2470-0029 (cit. on p. 68).

[163] J. Campbell, T. Neumann and Z. Sullivan,
Single-top-quark production in the t-channel at NNLO,
Journal of High Energy Physics **2021** (2021), ISSN: 1029-8479 (cit. on p. 68).

[164] E. Bothmann et al., *Event generation with Sherpa 2.2*, SciPost Physics **7** (2019),
ISSN: 2542-4653 (cit. on p. 68).

[165] T. Gleisberg and S. Höche, *Comix, a new matrix element generator*,
Journal of High Energy Physics **2008** (2008), ISSN: 1029-8479 (cit. on p. 68).

[166] F. Buccioni et al., *OpenLoops 2*, Eur. Phys. J. C **79** (2019) 866,
arXiv: 1907.13071 [hep-ph] (cit. on pp. 68, 69).

[167] S. Schumann and F. Krauss,
A parton shower algorithm based on Catani-Seymour dipole factorisation,
Journal of High Energy Physics **2008** (2008), ISSN: 1029-8479 (cit. on p. 68).

[168] S. Höche, F. Krauss, M. Schönherr and F. Siegert,
A critical appraisal of NLO+PS matching methods,
Journal of High Energy Physics **2012** (2012), ISSN: 1029-8479 (cit. on pp. 68, 69).

[169] S. Höche, F. Krauss, M. Schönherr and F. Siegert,
QCD matrix elements + parton showers. The NLO case,
Journal of High Energy Physics **2013** (2013), ISSN: 1029-8479 (cit. on pp. 68, 69).

[170] S. Catani, F. Krauss, B. R. Webber and R. Kuhn, *QCD Matrix Elements + Parton Showers*,
Journal of High Energy Physics **2001** (2001), ISSN: 1029-8479 (cit. on pp. 68, 69).

[171] S. Höche, F. Krauss, S. Schumann and F. Siegert,
QCD matrix elements and truncated showers, Journal of High Energy Physics **2009** (2009),
ISSN: 1029-8479 (cit. on pp. 68, 69).

[172] C. Anastasiou, L. Dixon, K. Melnikov and F. Petriello, *High-precision QCD at hadron colliders: Electroweak gauge boson rapidity distributions at next-to-next-to leading order*,
Physical Review D **69** (2004), ISSN: 1550-2368 (cit. on p. 68).

[173] LHC Higgs Cross Section Working Group,
Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector,
CERN Yellow Rep. Monogr. **2** (2017) 1, arXiv: 1610.07922 [hep-ph] (cit. on p. 69).

[174] H. B. Hartanto, B. Jäger, L. Reina and D. Wackeroth,
Higgs boson production in association with top quarks in the POWHEG BOX,
Physical Review D **91** (2015), ISSN: 1550-2368 (cit. on p. 69).

[175] T. Gleisberg and S. Hoeche, *Comix, a new matrix element generator*, JHEP **12** (2008) 039,
arXiv: 0808.3674 [hep-ph] (cit. on p. 69).

[176] S. Schumann and F. Krauss,
A Parton shower algorithm based on Catani-Seymour dipole factorisation,
JHEP **03** (2008) 038, arXiv: 0709.1027 [hep-ph] (cit. on p. 69).

[177] F. Cascioli, P. Maierhofer and S. Pozzorini, *Scattering Amplitudes with Open Loops*,
Phys. Rev. Lett. **108** (2012) 111601, arXiv: 1111.5206 [hep-ph] (cit. on p. 69).

[178] A. Denner, S. Dittmaier and L. Hofer,
Collier: a fortran-based Complex One-Loop LIbrary in Extended Regularizations,
Comput. Phys. Commun. **212** (2017) 220, arXiv: 1604.06792 [hep-ph] (cit. on p. 69).

[179] B. Fuks, K. Hagiwara, K. Ma and Y.-J. Zheng,
Signatures of toponium formation in LHC run 2 data, Phys. Rev. D **104** (2021) 034023,
arXiv: 2102.11281 [hep-ph] (cit. on pp. 69, 70).

[180] M. Jezabek, J. H. Kuhn and T. Teubner,
Momentum distributions in $t\bar{t}$ production and decay near threshold, Z. Phys. C **56** (1992) 653
(cit. on p. 69).

[181] A. D. Yaghjian, *Maxwell's derivation of the Lorentz force from Faraday's law*,
arXiv preprint arXiv:1911.04605 (2019) (cit. on p. 71).

[182] ATLAS Collaboration, *The ATLAS Collaboration Software and Firmware*, tech. rep.,
CERN, 2021, URL: https://cds.cern.ch/record/2767187 (cit. on p. 71).

[183] J. Pequenao and P. Schaffner,
"How ATLAS detects particles: diagram of particle paths in the detector", 2013,
URL: https://cds.cern.ch/record/1505342 (cit. on p. 73).

[184] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, Eur. Phys. J. C **83** (2023) 982, arXiv: 2212.09379 [hep-ex]
(cit. on p. 73).

[185] ATLAS Collaboration, *Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC*, The European Physical Journal C **77** (2017),
ISSN: 1434-6052 (cit. on p. 73).

[186] *ATLAS tracking software tutorial*,
https://atlassoftwaredocs.web.cern.ch/internal-links/tracking-tutorial/index.html, Accessed: 2025-10-13, 2025 (cit. on p. 74).

[187] ATLAS Collaboration, *Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2*, The European Physical Journal C **77** (2017), ISSN: 1434-6052,
URL: http://dx.doi.org/10.1140/epjc/s10052-017-5225-7 (cit. on p. 75).

[188] R. Frühwirth, *Application of Kalman filtering to track and vertex fitting*,
Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **262** (1987) 444, ISSN: 0168-9002,
URL: https://www.sciencedirect.com/science/article/pii/0168900287908874
(cit. on p. 75).

[189] ATLAS Collaboration,
A neural network clustering algorithm for the ATLAS silicon pixel detector,
Journal of Instrumentation **9** (2014) P09009,
URL: https://doi.org/10.1088/1748-0221/9/09/P09009 (cit. on p. 75).

[190] ATLAS Collaboration, *ATLAS pixel cluster splitting using Mixture Density Networks*,
tech. rep., CERN, 2019, URL: https://cds.cern.ch/record/2687968 (cit. on p. 76).

[191] T. G. Cornelissen et al., *The global χ^2 track fitter in ATLAS*,
Journal of Physics: Conference Series **119** (2008) 032013,
URL: https://doi.org/10.1088/1742-6596/119/3/032013 (cit. on p. 76).

[192] T. Cornelissen et al., *The new ATLAS track reconstruction (NEWT)*,
Journal of Physics: Conference Series **119** (2008) 032014 (cit. on p. 77).

[193] ATLAS Collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton–proton collision data*,
Journal of Instrumentation **14** (2019), ISSN: 1748-0221 (cit. on p. 78).

[194] ATLAS Collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data*,
Journal of Instrumentation **14** (2019) P12006 (cit. on pp. 79–81).

[195] ATLAS Collaboration,
Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1,
The European Physical Journal C **77** (2017), ISSN: 1434-6052 (cit. on pp. 78, 83).

[196] W. Lampl et al., *Calorimeter Clustering Algorithms*, tech. rep., CERN, 2008,
URL: https://cds.cern.ch/record/1099735 (cit. on p. 78).

[197] R. Frühwirth, *A Gaussian-mixture approximation of the Bethe–Heitler model of electron energy loss by bremsstrahlung*, Computer Physics Communications **154** (2003) 131,
ISSN: 0010-4655 (cit. on p. 79).

[198] ATLAS Collaboration, *Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13$ TeV*,
The European Physical Journal C **79** (2019), ISSN: 1434-6052 (cit. on pp. 80, 81).

[199] ATLAS Collaboration,
Performance of electron and photon triggers in ATLAS during LHC Run 2,
Eur. Phys. J. C **80** (2020) 47, arXiv: 1909.00761,
URL: https://cds.cern.ch/record/2688248 (cit. on p. 81).

[200] ATLAS Collaboration,
Electron and photon energy calibration with the ATLAS detector using LHC Run 1 data,
The European Physical Journal C **74** (2014), ISSN: 1434-6052,
URL: http://dx.doi.org/10.1140/epjc/s10052-014-3071-4 (cit. on p. 80).

[201] ATLAS Collaboration, *Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV*, The European Physical Journal C **81** (2021),
ISSN: 1434-6052, URL: http://dx.doi.org/10.1140/epjc/s10052-021-09233-2
(cit. on pp. 81, 82).

[202] ATLAS Collaboration, *Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s}$ =13 TeV*, The European Physical Journal C **76** (2016), ISSN: 1434-6052 (cit. on p. 81).

[203] ATLAS Collaboration, *Studies of the muon momentum calibration and performance of the ATLAS detector with pp collisions at $\sqrt{s}$ = 13 TeV*, The European Physical Journal C **83** (2023), ISSN: 1434-6052 (cit. on pp. 81, 83, 167).

[204] J. Illingworth and J. Kittler, *A survey of the hough transform*, Computer Vision, Graphics, and Image Processing **44** (1988) 87, ISSN: 0734-189X, URL: https://www.sciencedirect.com/science/article/pii/S0734189X88800331 (cit. on p. 81).

[205] ATLAS Collaboration, *Determination of jet calibration and energy resolution in proton–proton collisions at $\sqrt{s}$ = 8 TeV using the ATLAS detector*, The European Physical Journal C **80** (2020), ISSN: 1434-6052, URL: http://dx.doi.org/10.1140/epjc/s10052-020-08477-8 (cit. on p. 83).

[206] ATLAS Collaboration, *Jet reconstruction and performance using particle flow with the ATLAS Detector*, The European Physical Journal C **77** (2017), ISSN: 1434-6052 (cit. on p. 83).

[207] Y. Dokshitzer, G. Leder, S. Moretti and B. Webber, *Better jet clustering algorithms*, Journal of High Energy Physics **1997** (1997), ISSN: 1029-8479, URL: http://dx.doi.org/10.1088/1126-6708/1997/08/001 (cit. on pp. 84, 85).

[208] S. D. Ellis and D. E. Soper, *Successive combination jet algorithm for hadron collisions*, Physical Review D **48** (1993), ISSN: 0556-2821, URL: http://dx.doi.org/10.1103/PhysRevD.48.3160 (cit. on pp. 84, 85).

[209] M. Cacciari, G. P. Salam and G. Soyez, *The anti-ktjet clustering algorithm*, Journal of High Energy Physics **2008** (2008), ISSN: 1029-8479 (cit. on pp. 84, 85).

[210] ATLAS Collaboration, *Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s}$ = 8 TeV using the ATLAS detector*, The European Physical Journal C **76** (2016), ISSN: 1434-6052 (cit. on p. 86).

[211] A. Hoecker et al., *TMVA - Toolkit for Multivariate Data Analysis*, 2009, arXiv: physics/0703039 [physics.data-an] (cit. on p. 86).

[212] ATLAS Collaboration, *Performance of jet substructure techniques for large-R jets in proton-proton collisions at $\sqrt{s}$ = 7 TeV using the ATLAS detector*, JHEP **09** (2013) 076, arXiv: 1306.4945 [hep-ex] (cit. on p. 86).

[213] ATLAS Collaboration, *Jet reclustering and close-by effects in ATLAS run II*, (2017) (cit. on p. 87).

[214] ATLAS Collaboration, *Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s}$ = 13 TeV with the ATLAS detector*, 2021, arXiv: 2007.02645 [hep-ex] (cit. on pp. 88, 89, 119, 120).

[215] M. Cacciari, G. P. Salam and G. Soyez, *FastJet User Manual*, Eur. Phys. J. C **72** (2012) 1896, arXiv: 1111.6097 [hep-ph] (cit. on p. 88).

[216] ATLAS Collaboration,
Jet energy measurement with the ATLAS detector in proton–proton collisions at $\sqrt{s}$ = 7 TeV,
The European Physical Journal C **73** (2013), ISSN: 1434-6052 (cit. on p. 88).

[217] I. H. Witten, E. Frank, M. A. Hall and C. J. Pal, "Chapter 10 - Deep learning",
Data Mining (Fourth Edition), Fourth Edition, Morgan Kaufmann, 2017 417,
ISBN: 978-0-12-804291-5, URL:
https://www.sciencedirect.com/science/article/pii/B9780128042915000106
(cit. on p. 90).

[218] ATLAS Collaboration,
ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset,
The European Physical Journal C **83** (2023), ISSN: 1434-6052 (cit. on pp. 90–92).

[219] ATLAS Collaboration, *Configuration and performance of the ATLAS b-jet triggers in Run 2*,
The European Physical Journal C **81** (2021), ISSN: 1434-6052 (cit. on p. 91).

[220] ATLAS Collaboration,
Monte Carlo to Monte Carlo scale factors for flavour tagging efficiency calibration, tech. rep.,
CERN, 2020, URL: https://cds.cern.ch/record/2718610 (cit. on p. 92).

[221] ATLAS Collaboration, *Transforming jet flavour tagging at ATLAS*,
Nature Commun. **17** (2026) 541, arXiv: 2505.19689 [hep-ex] (cit. on pp. 92, 94).

[222] ATLAS Collaboration, *Graph Neural Network Jet Flavour Tagging with the ATLAS Detector*,
(2022) (cit. on p. 92).

[223] ATLAS Collaboration,
Deep Sets based Neural Networks for Impact Parameter Flavour Tagging in ATLAS, (2020)
(cit. on p. 93).

[224] ATLAS Collaboration, *Performance of missing transverse momentum reconstruction with the ATLAS detector using proton–proton collisions at* $\sqrt{s}$ *= 13 TeV*,
The European Physical Journal C **78** (2018), ISSN: 1434-6052 (cit. on pp. 93, 121).

[225] ATLAS Collaboration,
Performance of electron and photon triggers in ATLAS during LHC Run 2,
The European Physical Journal C **80** (2020), ISSN: 1434-6052 (cit. on p. 100).

[226] ATLAS Collaboration, *Performance of the ATLAS muon triggers in Run 2*,
Journal of Instrumentation **15** (2020), ISSN: 1748-0221 (cit. on p. 100).

[227] M. J. Fenton et al.,
Permutationless many-jet event reconstruction with symmetry preserving attention networks,
Phys. Rev. D **105** (11 2022) 112008 (cit. on pp. 101–103).

[228] A. Shmakov et al., *SPANet: Generalized permutationless set assignment for particle physics using symmetry preserving attention*, SciPost Phys. **12** (2022) 178 (cit. on p. 101).

[229] M. J. Fenton et al., *Reconstruction of unstable heavy particles using deep symmetry-preserving attention networks*, Communications Physics **7** (2024), ISSN: 2399-3650
(cit. on pp. 101, 102, 104).

[230] J. Erdmann et al.,
A likelihood-based reconstruction algorithm for top-quark pairs and the KLFitter framework,
Nucl. Instrum. Meth. A **748** (2014) 18, arXiv: 1312.5595 [hep-ex] (cit. on p. 102).

[231] J. Erdmann, T. Kallage, K. Kröninger and O. Nackenhorst,
From the bottom to the top—reconstruction of $t\bar{t}$ events with deep learning,
JINST **14** (2019) P11015, arXiv: 1907.11181 [hep-ex] (cit. on p. 102).

[232] L. Hertel, J. Collado, P. Sadowski, J. Ott and P. Baldi,
Sherpa: Robust hyperparameter optimization for machine learning,
SoftwareX **12** (2020) 100591, ISSN: 2352-7110, URL:
https://www.sciencedirect.com/science/article/pii/S2352711020303046
(cit. on p. 104).

[233] R. Ferraro and M. Thibeault,
Generic composition of boosts: an elementary derivation of the Wigner rotation,
European Journal of Physics **20** (1999) 143,
URL: https://doi.org/10.1088/0143-0807/20/3/003 (cit. on p. 106).

[234] M. Cacciari, G. P. Salam and G. Soyez, *The catchment area of jets*,
Journal of High Energy Physics **2008** (2008) 005,
URL: https://doi.org/10.1088/1126-6708/2008/04/005 (cit. on p. 107).

[235] ATLAS Collaboration, *Search for heavy neutral Higgs bosons decaying into a top quark pair in 140 fb^{-1} of proton-proton collision data at $\sqrt{s}$ = 13 TeV with the ATLAS detector*,
JHEP **08** (2024) 013, arXiv: 2404.18986 [hep-ex] (cit. on p. 107).

[236] T. Reisch, *Disentangling Quantum Mechanics in Top Quark Pairs at the ATLAS Experiment, CERN-THESIS-2024-341*, 2024, URL: https://cds.cern.ch/record/2925172
(cit. on p. 110).

[237] G. D'Agostini, *A multidimensional unfolding method based on Bayes' theorem*,
Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **362** (1995) 487, ISSN: 0168-9002
(cit. on p. 112).

[238] T. Adye, "Unfolding algorithms and tests using RooUnfold", *PHYSTAT 2011*,
Geneva: CERN, 2011 313, arXiv: 1105.1160 [physics.data-an] (cit. on p. 112).

[239] *TREXFitter Profile Likelihood Unfolding*, https://trexfitter-docs.web.cern.ch/trexfitter-docs/latest/tutorials/unfolding/,
Accessed: 2026-04-06, 2026 (cit. on pp. 112, 149, 156).

[240] S. Schmitt, *TUnfold: an algorithm for correcting migration effects in high energy physics*,
JINST **7** (2012) T10003, arXiv: 1205.6201 [physics.data-an] (cit. on p. 112).

[241] ATLAS Collaboration,
Tools for estimating fake/non-prompt lepton backgrounds with the ATLAS detector at the LHC,
JINST **18** (2023) T11004, arXiv: 2211.16178 [hep-ex] (cit. on pp. 115, 117).

[242] M. Cacciari, M. Czakon, M. Mangano, A. Mitov and P. Nason, *Top-pair production at hadron colliders with next-to-next-to-leading logarithmic soft-gluon resummation*,
Phys. Lett. B **710** (2012) 612, arXiv: 1111.5869 [hep-ph] (cit. on p. 121).

[243] J. H. Kühn, A. Scharf and P. Uwer,
Weak Interactions in Top-Quark Pair Production at Hadron Colliders: An Update,
Phys. Rev. D **91** (2015) 014020, arXiv: 1305.5773 [hep-ph] (cit. on p. 122).

[244] ATLAS Collaboration, *Measurement of the top quark mass in the $t\bar{t} \to$ lepton+jets channel from $\sqrt{s} = 8$ TeV ATLAS data and combination with previous results*,
Eur. Phys. J. C **79** (2019) 290, arXiv: 1810.01772 [hep-ex] (cit. on p. 122).

[245] T. Ježo, J. M. Lindert, P. Nason, C. Oleari and S. Pozzorini, *An NLO+PS generator for $t\bar{t}$ and Wt production and decay including non-resonant and interference effects*,
Eur. Phys. J. C **76** (2016) 691, arXiv: 1607.04538 [hep-ph] (cit. on pp. 122, 180).

[246] L. A. Harland-Lang, A. D. Martin, P. Motylinski and R. S. Thorne,
Parton distributions in the LHC era: MMHT 2014 PDFs, Eur. Phys. J. C **75** (2015) 204,
arXiv: 1412.3989 [hep-ph] (cit. on p. 122).

[247] ATLAS Collaboration, *Studies on the improvement of the matching uncertainty definition in top-quark processes simulated with Powheg+Pythia 8*, tech. rep., CERN, 2023,
URL: https://cds.cern.ch/record/2872787 (cit. on p. 123).

[248] H. Brooks and P. Skands, *Coherent showers in decays of colored resonances*,
Phys. Rev. D **100** (2019) 076006, arXiv: 1907.08980 [hep-ph] (cit. on p. 123).

[249] ATLAS Collaboration, *A study of different colour reconnection settings for Pythia8 generator using underlying event observables*, tech. rep., CERN, 2017,
URL: https://cds.cern.ch/record/2262253 (cit. on p. 124).

[250] ATLAS Collaboration, *Measurement of $t\bar{t}$ production in association with additional b-jets in the $e\mu$ final state in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*,
JHEP **01** (2025) 068, arXiv: 2407.13473 [hep-ex] (cit. on p. 125).

[251] ATLAS Collaboration, *Measurement of top-quark pair production in association with charm quarks in proton–proton collisions at s=13 TeV with the ATLAS detector*,
Phys. Lett. B **860** (2025) 139177, arXiv: 2409.11305 [hep-ex] (cit. on p. 125).

[252] J. Alwall et al., *Comparative study of various algorithms for the merging of parton showers and matrix elements in hadronic collisions*, Eur. Phys. J. C **53** (2008) 473,
arXiv: 0706.2569 [hep-ph] (cit. on p. 125).

[253] ATLAS Collaboration,
Search for the standard model Higgs boson produced in association with top quarks and decaying into a $b\bar{b}$ pair in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector,
Phys. Rev. D **97** (2018) 072016, arXiv: 1712.08895 [hep-ex] (cit. on p. 125).

[254] R. Frederix and S. Frixione, *Merging meets matching in MC@NLO*, JHEP **12** (2012) 061,
arXiv: 1209.6215 [hep-ph] (cit. on p. 140).

[255] K. Cranmer, G. Lewis, L. Moneta, A. Shibata and W. Verkerke,
HistFactory: A tool for creating statistical models for use with RooFit and RooStats, tech. rep.,
New York U., 2012, URL: https://cds.cern.ch/record/1456844 (cit. on pp. 150–152).

[256] ATLAS Collaboration, *AtlFast3: The Next Generation of Fast Simulation in ATLAS*,
Comput. Softw. Big Sci. **6** (2022) 7, arXiv: 2109.02551 [hep-ex] (cit. on p. 158).

[257] *Description of the different smoothing algorithms- Tomas Dado*, https://gitlab.cern.ch/atlas/statistics/SystematicSmoothingTool/-/blob/master/doc/algorithm_description.md, Accessed: 2026-04-29 (cit. on p. 158).

[258] *TREXFitter- nuisance parameter ranking*, https://trexfitter-docs.web.cern.ch/trexfitter-docs/latest/inference/ranking/, Accessed: 2026-04-29 (cit. on p. 171).

[259] A. Pinto et al., *Uncertainty components in profile likelihood fits*, Eur. Phys. J. C **84** (2024) 593, arXiv: 2307.04007 [physics.data-an] (cit. on p. 171).

[260] M. V. Garzelli, G. Limatola, S.-O. Moch, M. Steinhauser and O. Zenaiev, *Threshold Top-Quark Pair-Production: Cross Sections and Key Uncertainties*, (2026), arXiv: 2604.09485 [hep-ph] (cit. on p. 180).

[261] B. Fuks, A. Hossain and J. Keaveney, *Statistical indications of toponium formation in top quark pair production*, Physics Letters B **873** (2026) 140179, ISSN: 0370-2693, URL: http://dx.doi.org/10.1016/j.physletb.2026.140179 (cit. on p. 180).

[262] J. Mazzitelli et al., *Next-to-Next-to-Leading Order Event Generation for Top-Quark Pair Production*, Physical Review Letters **127** (2021), ISSN: 1079-7114, URL: http://dx.doi.org/10.1103/PhysRevLett.127.062001 (cit. on p. 180).

[263] Y. Afik and J. R. M. de Nova, *Quantum Discord and Steering in Top Quarks at the LHC*, Phys. Rev. Lett. **130** (2023) 221801, arXiv: 2209.03969 [quant-ph] (cit. on p. 180).

[264] D. Atwood, S. Bar-Shalom, G. Eilam and A. Soni, *CP violation in top physics*, Phys. Rept. **347** (2001) 1, arXiv: hep-ph/0006032 (cit. on p. 181).

[265] C. D. White and M. J. White, *The magic of entangled top quarks*, 2024, arXiv: 2406.07321 [hep-ph], URL: https://arxiv.org/abs/2406.07321 (cit. on p. 181).

[266] CMS Collaboration, *Observation of magic states of top quark pairs produced in proton-proton collisions at* $\sqrt{s}$ = 13 TeV, tech. rep., CERN, 2025, URL: https://cds.cern.ch/record/2926751 (cit. on p. 181).

[267] R. Aoude et al., *Radiation effects on the entanglement of fermion pairs at colliders*, 2026, arXiv: 2604.16268 [hep-ph], URL: https://arxiv.org/abs/2604.16268 (cit. on p. 181).

[268] CMS Collaboration, *Study of spin correlations in Higgs boson decays to four leptons at CMS*, tech. rep., CERN, 2025, URL: https://cds.cern.ch/record/2949823 (cit. on p. 181).

[269] Y. Afik, Y. Kats, J. R. M. de Nova, A. Soffer and D. Uzan, *Entanglement and Bell nonlocality with bottom-quark pairs at hadron colliders*, Phys. Rev. D **111** (2025) L111902, arXiv: 2406.04402 [hep-ph] (cit. on p. 181).

[270] Y. Afik et al., *Quantum information meets high-energy physics: input to the update of the European strategy for particle physics*, Eur. Phys. J. Plus **140** (2025) 855, arXiv: 2504.00086 [hep-ph] (cit. on p. 181).

List of Figures

2.1 A schematic diagram of the elementary particles in the SM. The different colours indicate the fermion and boson sectors as discussed in the text. For each particle, its mass, electric charge, and spin are shown [32]. The values of the particle properties, shown in the diagram, are taken from Ref. [31]. 9

2.2 Illustration of the steps in the hadronisation process. (i) The initial quarks, q_0 and $\bar{q}_0$, begin to move apart. (ii) A narrow flux tube forms between them. (iii) The string subsequently breaks to create a new quark–antiquark pair. (iv) Repeated string breaking occurs until (v) all quarks and antiquarks have sufficiently low energy to combine into colourless hadrons. The hadronisation process results in two jets of hadrons. [35]. 12

2.3 An illustration of the Higgs potential in lower dimensions, where ϕ is taken to be a single complex scalar field rather than an $SU(2)$ doublet. The parameters satisfy $\mu^2 < 0$ and $\lambda > 0$, leading to the characteristic "Mexican hat" shape of the potential [41]. 16

3.1 PDF4LHC PDF sets [65] as a function of x, the proton momentum fraction, for different choices of the factorisation scale, denoted by μ in the figures [66]. The left plot corresponds to a factorisation scale of 10^2 GeV^2, while the right one corresponds 10^4 GeV^2. The y-axis shows the quantity $xf(x,\mu^2)$, such that the distributions represent the fraction of the proton momentum carried by different partons. 24

3.2 LO Feynman diagrams for $t\bar{t}$ production: (a) quark–antiquark annihilation and (b–d) gluon fusion via s-, t-, and u-channels. 24

3.3 Feynman diagrams illustrating the three different $t\bar{t}$ decay channels. 25

3.4 Differential cross-section for $t\bar{t}$ production as a function of the $t\bar{t}$ invariant mass at the LHC with $\sqrt{s} = 14$ TeV, separated by left-handed and right-handed helicity configurations for both production modes [69]. 28

3.5 Illustration of a top-quark decay, in the top-quark rest frame, into a b quark and a W^+ boson with (a) right-handed helicity (b) left-handed helicity and (c) longitudinal (zero) helicity [71, 72]. The single arrows indicate the directions of the particle momenta, while the double arrows denote the spin directions of the top quark (lower arrow) and of the b-quark and W boson (upper arrows). 30

3.6 The top-quark decay system as viewed in the W-boson rest frame (left), and boosted back into the top-quark rest frame (right). The helicity angle, $\theta_{W\,\mathrm{hel}}$, is defined as the direction of the down-type quark with respect to $-\hat{b}$ [73]. 31

3.7 Spin analysing powers as a function of $|c_{W\text{hel}}|$ for the optimal hadronic polarimeter q_{opt} (black), soft quark q_{soft} (red), and b-quark (blue) [73]. 32
3.8 Two different orthonormal bases in the CM frame: helicity basis on the left and beam basis on the right [75]. 34
3.9 Measure of entanglement, termed concurrence, as a function of the invariant mass of the $t\bar{t}$ system, $m_{t\bar{t}}$ (denoted by $M_{t\bar{t}}$ in these figures), and the production angle Θ in the $t\bar{t}$ CM frame. Black lines represent the boundaries between separability and entanglement [75]. 36
3.10 Illustration of the time evolution of a non-relativistic $t\bar{t}$ system leading to the formation of a quasi-bound-state via multiple-gluon exchange. The initial plane-wave components describing the $t\bar{t}$ system propagate outwards from the origin x=0. These waves evolve until they reach the effective potential barrier at a characteristic distance given by the Bohr radius, $a_0 = (\alpha_s m_t)^{-1}$. At times $t \sim a_0/\alpha_s$, the typical waves related to the bound states reach the barrier, and the formation of resonance-like quasi-bound configurations becomes relevant. The system remains in this state until the t or $\bar{t}$ decays at the timescale Γ_Θ^{-1}, the lifetime of this bound state [20]. 42
3.11 Distributions of the $t\bar{t}$ invariant mass for $gg \to t\bar{t}$ in the color-singlet (red), and color-octet (blue) states, and the sum of them (black). The solid lines represent the $J = 0$ contribution and dashed lines are calculated as the sum of contributions from all J values [90]. 42

4.1 A schematic illustration of the LHC complex. The four large-scale detectors are shown at the four main IPs, together with the pre-accelerator chain (BOOSTER, PS, and SPS), which boosts the protons before injecting them into the main LHC ring [97]. 46
4.2 (a) Integrated luminosity as a function of time delivered to ATLAS (green), recorded by ATLAS (yellow), and passing data-quality criteria (blue) [106] (b) Distribution of the mean number of interactions per bunch crossing, $\langle\mu\rangle$, at the ATLAS experiment for pp collisions during Run-2. The quantity $\langle\mu\rangle$ is calculated from the instantaneous luminosity assuming an inelastic pp collision cross-section of 80 mb at $\sqrt{s} = 13$ TeV [107]. 49
4.3 Cut-away view of the ATLAS detector [108], shown in its Run-3 configuration, which differs from the Run-2 setup primarily through upgrades to the Small Wheels (Section 4.2.4). 50
4.4 An illustration of the ATLAS coordinate system. Image taken from Ref. [109] and modified. 51
4.5 Transverse section of the ATLAS ID. Going outward from the beamline, the ID subsystems are: the Silicon Pixel Detector, the Semiconductor Tracker (SCT), and the Transition-Radiation Tracker (TRT) [108]. 53
4.6 Illustration of the ATLAS calorimeter system. The electromagnetic calorimeter surrounds the ID and is overlaid by the hadronic calorimeter. The main components are labelled in the figure [108]. 54
4.7 (a) Illustration of an electromagnetic shower initiated by a high-energy photon. (b) Schematic depiction of a hadronic shower showing its electromagnetic and hadronic components [113]. 55
4.8 Sketch of the structure of an ATLAS electromagnetic LAr calorimeter barrel module [98]. 56

4.9 A schematic view of the ATLAS magnet system: the solenoid (blue), the barrel toroid (pink) and the two end-cap toroids (green) [116]. 57

4.10 Illustration of the ATLAS Muon Spectrometer in Run-3, showing the new small wheels (NSW) that replaced the innermost end-cap wheels from Run-2 [108]. 58

4.11 An illustration of the future ATLAS ITk, showing the pixel detector (in green), the strip detector (in blue), and the high-granularity timing detector (HGTD, in orange) [119]. 60

5.1 Illustration of a pp collision event corresponding to the process $gg \to t\bar{t} \to q'\bar{q}b\,\mu\nu_\mu\bar{b}$, as modelled by an event generator. The incoming partons are depicted as dark-green ovals, while the hard-scattering interaction is represented by the large red filled circle. The produced top quarks are indicated by smaller red filled circles. The subsequent parton shower evolution is shown in blue, and additional soft interactions from the underlying event are illustrated in purple. Final-state partons undergo hadronisation (light-green filled circles), after which the resulting hadrons decay (dark-green filled circles). Photon radiation, shown in yellow, may occur at various stages of the event evolution. This schematic is reproduced from Ref. [133]. 63

5.2 Illustration of the different levels of event description for the $t\bar{t}$ lepton+jets resolved decay channel: parton level (left, blue), particle level (middle, green), and detector level (right, red). At the parton level, the decay products are represented by quarks and leptons before hadronisation. At the particle and detector levels, the quarks have hadronised into particles that are clustered into jets, illustrated by the conical jet shapes in the diagram. The faded ATLAS detector in the first two panels indicates that detector effects are not included at the parton or particle levels. 66

5.3 Feynman diagrams of single-top production in the Wt channel beyond LO. Top and antitop quarks are represented with double lines. 68

6.1 Schematic illustration of a sector of the cross-section of the ATLAS detector, showing the trajectories of different particles as they traverse the various detector components. Dashed lines indicate particles that escape detection. Figure adapted from Ref. [183]. 73

6.2 Schematic illustration of the perigee representation of the track parameters, where the perigee corresponds to the point of closest approach of the trajectory to the reference point [186]. 74

6.3 Different cluster topologies formed on a pixel sensor of a pixel detector layer. Arrows indicate particle tracks, and colours correspond to energy deposits from individual particles [187]. 75

6.4 Illustration of the ATLAS ID track reconstruction procedure. Space points are shown in yellow, while seeds formed from combinations of space points are indicated in blue. The dashed blue seed illustrates a situation where two seeds correspond to the trajectory of the same charged particle. A seed rejected based on the requirement that it needs to be consistent with a particle originating from the nominal interaction point is shown in green. The black track represents a track candidate reconstructed using CKF. The red tracks correspond to the fully reconstructed silicon tracks obtained after the global χ^2 fit, where an extension into the TRT may also be performed [26]. . . . 76

6.5 Results of f_{lost} as a function of jet p_T in Run-3 data collected in 2022 and the corresponding MC simulation. The measurement is based on dE/dx distributions in the range $0.6 < dE/dx < 3.0$ MeV g^{-1} cm^2. The error bands include both statistical and systematic uncertainties [26]. 78

6.6 Illustration of the trajectory of an electron in the ATLAS detector. The red curve represents the path of the electron before it enters the electromagnetic calorimeter, while the dashed line indicates the trajectory of a photon radiated by the electron due to its interaction with the tracker material [194]. 79

6.7 Illustrative examples of reconstructed jet shapes obtained with different jet clustering algorithms in the $y - \phi$ plane. The z-axis denotes the p_T of the reconstructed jet. For simplicity, in all cases the radius parameter is set to 1.0. Different colours correspond to different jets. [209]. 85

6.8 Diagram illustrating the topology of a $t\bar{t}$ event in the lepton+jets decay channel containing a boosted top quark. The hadronically decaying top quark is reconstructed as a RC large-R jet (shown in pink), which contains two small-R jets (blue) originating from the hadronic decay of the W boson and a small-R jet identified as a b-jet via b-tagging (green). 87

6.9 Illustration of the secondary vertex, decay length (L_{xy}), and the transverse impact parameters d_0 for a b-hadron decay within a jet emanating from the primary interaction vertex [219]. 91

6.10 Distribution of the DL1r b-jet discriminant output for b-jets (solid), c-jets (dashed-dotted), and light-flavour jets (dashed). The curves represent the normalised probability distributions of the DL1r b-jet discriminant output for each jet flavour, i.e. the probability of a jet of a given flavour to obtain a particular b-tagging discriminant value. The distribution peaks at higher values for b-jets, while c- and light-flavour jets are concentrated at lower values, with an intermediate enhancement region visible for c-jets [218]. 91

6.11 Rejection rates for c-jet (solid), light-jet (dash–dotted), and τ-jet (dashed) as a function of the b-jet tagging efficiency for jets in simulated $t\bar{t}$ events with $20 < p_T < 250$ GeV. Results are shown for the GN2 (light blue) and DL1d (dark orange) taggers. The lower panels display the relative performance of GN2 with respect to DL1d. Figure taken from Ref. [221]. 94

7.1 Diagram of the SPANet architecture. Arranged from left to right, the inputs are denoted by ϵ_i, the assignment outputs by $\mathcal{P}_1 \ldots \mathcal{P}_m$, the regression outputs by η_ν and $m_{t\bar{t}}$, and the classification output by $\mathcal{S}/\mathcal{B}$ [229]. In this measurement, $m = 2$, corresponding to the two top-quark assignment probability distributions. 102

7.2 Fraction of events with correctly matched, incorrectly matched, and non-reconstructable down-type jets as a function of the `spanet_hadtop_detection` score. The fraction of correctly matched down-type jets clearly increases with higher values of this score [236]. 110

7.3 Schematic representation of the calibration curve. The nominal SM prediction obtained from simulation is shown as a red triangle, while alternative entanglement scenarios are represented by black square points derived through reweighting. These points are linearly interpolated to construct the calibration curve (yellow line). Statistical and systematic variations are evaluated at each point, resulting in additional calibration curves that define the overall uncertainty band shown in grey. A possible measured value of D is indicated by the arrow and is mapped to the corresponding particle-level value (illustrated by dashed lines) using the calibration curve. The associated uncertainty on this mapping is shown by the vertical uncertainty bands. 113

7.4 Reweighted $\cos\varphi$ distributions at the parton level, obtained using different χ values in Equation 7.3, used to construct the calibration curve. 114

7.5 Fake efficiencies estimated in the control region for electrons (left) and muons (right) as a function of the lepton transverse momentum, p_{T}, and pseudorapidity, η. 117

7.6 Distribution of the invariant mass of the $t\bar{t}$ system, $m_{t\bar{t}}$, at the detector level, showing the breakdown of signal and background contributions. The ratio between data and prediction is shown in the lower panel. The uncertainties include statistical (dashed red) and all systematic contributions considered in this analysis. No selections on $\beta_{t\bar{t}}$ and SPANet performance variables are applied. The y-axis is in the logarithmic scale. 126

7.7 Distribution of the angular variable $\cos\varphi$ at the detector level in the signal region ($m_{t\bar{t}} < 390\,\mathrm{GeV}$), showing the breakdown of signal and background contributions. The ratio between data and prediction is shown in the lower panel. The uncertainties include statistical (dashed red) and all systematic contributions considered in this analysis. 127

7.8 Detector-level distributions of the transverse momentum, p_T, of the lepton and the down-type jet, which serve as spin analysers used in the construction of the angular observable $\cos\varphi$, in the signal region. The ratio between data and prediction is shown in the lower panel. The uncertainties include statistical (dashed red) and all systematic contributions considered in this analysis 128

7.9 Distributions of the angular variable $\cos\varphi$ at the detector level in the validation regions, showing the breakdown of signal and background contributions. The ratio between data and prediction is shown in the lower panel. The uncertainties include statistical (dashed red) and all systematic contributions considered in this analysis. . . 128

7.10 Calibration curve in the signal region ($m_{t\bar{t}} < 390\,\mathrm{GeV}$). The black filled circular marker represents the data while the SM expectation is shown by the red triangle marker. The grey square markers represent the alternative entanglement scenarios introduced in Section 7.7. 130

7.11 Calibration curves in the validation regions, similar to Figure 7.10. Note that the axes scales differ. 131

7.12 Calibration curve mapping the parton-level D value to the particle-level value, including statistical and modelling systematic uncertainties (except the PS and hadronisation uncertainty). The entanglement limit $D = -1/3$ propagated from parton to particle level, denoted by $D^{-1/3}_{\mathrm{particle}}$, is marked with a star. 133

7.13 Calibration curves mapping between different levels of event description for two generator setups: nominal POWHEG+PYTHIA (red) and alternative POWHEG+HERWIG (blue), used to estimate the PS and hadronisation uncertainty. A larger difference between the calibration curves is observed when mapping parton-level D to particle-level D. 135

7.14 Overview of the values of the entanglement observable D at the particle level in the signal and the two validation regions. Filled markers correspond to the nominal fit model, which include quasi-bound-state effects, while open markers represent the scenario without quasi-bound-state effects. The observed values, extracted using the calibration curve of POWHEG+PYTHIA with or without $t\bar{t}_{\mathrm{GFRW}}$ (without denoted by $\not{t\bar{t}}_{\mathrm{GFRW}}$), are shown as black circles. The expected values from POWHEG+PYTHIA and POWHEG+HERWIG are indicated by red triangles and blue squares, respectively. For the data points, the inner error bars denote the statistical uncertainties, while the outer error bars represent the total uncertainties. For the POWHEG+PYTHIA prediction (red triangles), the corresponding uncertainties are indicated by inner error bars and a shaded band, respectively. The entanglement limits, propagated from the parton to the particle level for both generators, are also shown together with their associated uncertainties. 136

7.15 Overview of the ATLAS dilepton entanglement analysis result. The particle-level D results in the signal and validation regions are compared for various MC models. The entanglement limit, converted from the parton-level value $D = -1/3$ to the particle level, is shown [17]. 139

7.16 Summary of the CMS dilepton entanglement measurement. Shown are the observed parton-level D values (black points) compared with MC predictions with (solid) and without (dashed) the η_t state. Inner (outer) error bars show the statistical (total) uncertainty in data. Light (dark) shaded regions indicate the MC statistical (total) uncertainty. The entanglement limit $D = -1/3$ is shown as a shaded band [23] 140

7.17 Measurement results of the entanglement observable D by CMS in the lepton+jets threshold region. The observed values are shown by black filled circles with the inner (outer) error bars showing statistical (total) uncertainty. The observed values are compared to POWHEG+PYTHIA, POWHEG+PYTHIA+η_t, POWHEG+HERWIG, MADGRAPH5_AMC@NLO+PYTHIA, and MINNLO+PYTHIA predictions. Observed (expected) significance relative to separable states (green band) is given in units of standard deviations σ. The boundary of the green band corresponds to the entanglement threshold $D = -1/3$, below which the $t\bar{t}$ system is entangled [24]. 141

8.1 Illustration of the region definitions used for unfolding at the reconstruction level (top, shown in red) and at the parton level (bottom, shown in blue). All regions at both levels are defined with the selections $m_{t\bar{t}} > 800$ GeV and $|\cos\Theta| < 0.4$. 149

8.2 Schematic illustration of the folding procedure from the parton level to the reconstruction level using the response matrix. In this example, both the parton- and reconstruction-level distributions are divided into 10 bins, resulting in a 10×10 response matrix. The matrix is predominantly diagonal, indicating that most events are reconstructed in the same bin as defined at the parton level. For instance, the third parton-level bin (highlighted in red) maps primarily onto the third bin of the folded reconstruction-level distribution, with smaller contributions to neighbouring bins arising from bin-to-bin migrations. 155
8.3 Migration matrix, acceptance, efficiency, and the resulting response matrix in SR1 e+jets. The z-axis of the migration matrix denotes the probability that an event reconstructed in a given reconstruction-level bin originates from a particular parton-level bin. These values are explicitly shown in the matrix itself. For the response matrix, the corresponding axis represents the same probability after accounting for detector acceptance and efficiency effects. 160
8.4 Migration matrix, acceptance, efficiency, and the resulting response matrix in SR2 e+jets. The z-axis of the migration matrix denotes the probability that an event reconstructed in a given reconstruction-level bin originates from a particular parton-level bin. These values are explicitly shown in the matrix itself. For the response matrix, the corresponding axis represents the same probability after accounting for detector acceptance and efficiency effects. 162
8.5 Pre- and post-fit distributions of $\cos\theta'$ in SR1, showing the MC signal and background contributions. The data used in the fit correspond to Asimov data. Hence, no data points are displayed. The lower panel shows the ratio of the data to the total MC prediction. The uncertainty band includes the full set of statistical and systematic uncertainties considered in the fit. 164
8.6 Pre- and post-fit distributions of $\cos\theta'$ in SR2, showing the MC signal and background contributions. The data used in the fit correspond to Asimov data. Hence, no data points are displayed. The lower panel shows the ratio of the data to the total MC prediction. The uncertainty band includes the full set of statistical and systematic uncertainties considered in the fit. 165
8.7 Pulls and constraints for NPs associated with $t\bar{t}$ modelling uncertainties. Black points indicate pulls, while horizontal bars represent constraints. The green and yellow bands correspond to the 1σ and 2σ ranges. 167
8.8 Pulls and constraints for NPs corresponding to the JER systematic uncertainties. . . . 167
8.9 Pulls and constraints for NPs corresponding to the JES systematic uncertainties. . . . 168
8.10 Pulls and constraints for NPs corresponding to the systematic uncertainties associated with electrons. 168
8.11 Pulls and constraints for NPs corresponding to the systematic uncertainties associated with muons. The systematic components of the reconstruction, isolation, and TTVA efficiencies are shown separately, as explained in the text. 169
8.12 Pulls and constraints for NPs corresponding to the flavour-tagging systematic uncertainties. FT EFF denotes flavour tagging efficiency uncertainties. The components labelled Eigen B, Eigen C, and Eigen Light correspond to the eigenvector variations associated with b-, c-, and light-flavour jets, respectively. The corresponding high-p_T extrapolation uncertainties are denoted by pT b Eigen, pT c Eigen, and pT light Eigen. 170

8.13 Pulls and constraints for NPs of experimental uncertainties related to pile-up reweighting, luminosity, and JVT efficiency. 170
8.14 Correlation matrix for the POIs and NPs in the fit. The values represent the correlation coefficients between pairs of parameters, expressed in percentage. Only parameters with at least one correlation exceeding 20% in magnitude are shown. 172
8.15 Post-fit ranking plot showing the impact of the 15 most significant uncertainty contributions to the extracted value of the POI, μ. The post-fit up ($\sigma_{\mathrm{POI}}^{\mathrm{up}}\,\rho\,\sigma_{\mathrm{NP}}^{\mathrm{up}}$) and down ($\sigma_{\mathrm{POI}}^{\mathrm{down}}\,\rho\,\sigma_{\mathrm{NP}}^{\mathrm{down}}$) impacts are shown. Pulls and constraints on the corresponding NPs are also indicated (black points with horizontal bars), with the parameter axis shown at the bottom. 173
8.16 Shape of the JET Flavour Response systematic for a given parton level bin of the $\cos\theta'$ distribution of the $t\bar{t}$ signal folded to the reconstruction level in SR1, μ +jets channel. The up variation is given by the orange lines ($+1\sigma$) and the down variation is given by the violet lines (-1σ). The dashed lines correspond to the raw systematic distributions and the solid lines represent the smoothed and symmetrised distributions. 174
8.17 Shape of the JET Flavour Response systematic for a given parton level bin of the $\cos\theta'$ distribution of the $t\bar{t}$ signal folded to the reconstruction level in SR2, μ +jets channel. The up variation is given by the orange lines ($+1\sigma$) and the down variation is given by the violet lines (-1σ). The dashed lines correspond to the raw systematic distributions and the solid lines represent the smoothed and symmetrised distributions. 175
8.18 Measurement of $\tilde{D}$ (corresponding to D_3 in this thesis) by CMS in the lepton+jets boosted regime. The points show statistical (inner) and total (outer) uncertainties, compared to predictions from Powheg+Pythia, Powheg+Herwig, MadGraph5_aMC@NLO+Pythia, and MiNNLO+Pythia. The observed (expected) significance relative to separable states (green band) is quoted in units of σ. The boundary of the green band corresponds to the entanglement threshold $\tilde{D} = 1/3$, above which the $t\bar{t}$ system is entangled [24]. 176

9.1 Schematic illustration of the relation between the different quantum correlations. . . 181

A.1 Fits of various kinematic parameters entering the χ^2 algorithm. 184

C.1 Migration matrix, acceptance, efficiency, and the resulting response matrix SR1 μ+jets. The z-axis of the migration matrix denotes the probability that an event reconstructed in a given reconstruction-level bin originates from a particular parton-level bin. These values are explicitly shown in the matrix itself. For the response matrix, the corresponding axis represents the same probability after accounting for detector acceptance and efficiency effects. 188
C.2 Migration matrix, acceptance, efficiency, and the resulting response matrix SR2 μ+jets. The z-axis of the migration matrix denotes the probability that an event reconstructed in a given reconstruction-level bin originates from a particular parton-level bin. These values are explicitly shown in the matrix itself. For the response matrix, the corresponding axis represents the same probability after accounting for detector acceptance and efficiency effects. 189

List of Tables

3.1 Branching fractions of the different $t\bar{t}$ decay channels, as reported in [31]. Here, all flavours of leptons are included $l = e, \mu, \tau$. 26
3.2 Spin analysing powers of potential top-quark spin analysers [70, 74]. 33

6.1 Definition of the electron isolations WPs relevant for this thesis. The first one is used in the threshold entanglement measurement, while the second one is used in the boosted analysis. Fixed-cone and variable-cone requirements are applied to the calorimeter- and track-based isolation variables, respectively [194, 198, 199]. 81
6.2 Definition of the muon isolation WPs used in this thesis. The first one is used in the threshold entanglement measurement, while the second one is used in the boosted analysis. 82
6.3 Table illustrating the criteria used for the various steps of the overlap removal procedure. 95

7.1 Selection criteria for electrons and muons. The looser criteria applied in the definition of the "fake" control regions are shown in italics and square brackets. 99
7.2 Selection criteria for small-R jets. 99
7.3 Summary of the single-lepton triggers used in this analysis and their requirements. . 100
7.4 Summary of the event selection criteria. 101
7.5 Summary of the fraction of correctly assigned jets for different selection requirements on the value of $\log_{10}(\chi^2)$, following the application of the χ^2 minimisation technique to determine the optimal combination of jets at the particle level. 108
7.6 Fraction of events with correctly matched, incorrectly matched or non-reconstructable down-type jets, for cases with and without a lower bound on the `spanet_hadtop_detection` score. 110
7.7 Definition of loose and tight leptons used in the non-prompt and fake lepton estimate 116
7.8 Event yield summary for the signal and validation regions. The quoted uncertainties on the predicted event yields include both statistical and systematic contributions, summed in quadrature. 129
7.9 A summary of the grouped sources of systematic uncertainty in terms of their absolute and relative impact on the observed and expected values of D after calibration from detector to particle level. 132

7.10 Ranking of the ten largest sources of systematic uncertainty in terms of their absolute and relative impact on the observed value of D after calibration from detector to particle level. NP-(X) denotes the X-th component within the set of jet or flavour-tagging uncertainties, as introduced in Section 7.9.1. The Jet JER Data vs. MC uncertainty corresponds to the component associated with differences in JER modelling between data and simulation. . . . 132
7.11 Ranking of the impact of the ten most dominant sources of systematic uncertainty on the calibration of the entanglement limit $D = -1/3$ from the parton to the particle level in terms of their absolute and relative size. Only the $t\bar{t}$ modelling uncertainties (except PS and hadronisation) are considered. . . . 134
7.12 Ranking of the largest sources of systematic uncertainty in terms of their absolute and relative impact on the observed value of D after calibration from detector to particle level, using the signal model that employs the simplified model, η_t, for quasi-bound-state effects. Impact of the η_t modelling uncertainties are also listed for comparison. . . . 138

8.1 Selection criteria for electrons and muons. . . . 145
8.2 Selection criteria for small-R and RC jets. . . . 146
8.3 Summary of the event selection criteria (Boosted analysis). . . . 147
8.4 Expected number of events per signal region prior to the fit, showing both the total expected yields and yields per process. The uncertainties include both statistical and systematic contributions. . . . 161
8.5 Summary of the POIs extracted from the PLU fit. The quoted uncertainties are decomposed into statistical and systematic components, listed in that order. . . . 163
8.6 Summary of the POIs extracted from the PLU fit for the different signal regions and their combination. The quoted uncertainties are decomposed into statistical and systematic components, listed in that order. . . . 166

B.1 Summary of results obtained using different selection cuts to optimise the entanglement signal region. . . . 185