Computationally Affordable Methods towards High-Dimensional Molecular Systems

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

vorgelegt von
Álvaro Fernández Corral

Hamburg
2025

Gutachter/innen der Dissertation: Prof. Dr. Jochen Küpper
Prof. Dr. Armin Iske

Zusammensetzung der Prüfungskommission: Prof. Dr. Jochen Küpper
Prof. Dr. Armin Iske
Dr. Andrey Yachmenev
Prof. Dr. Robin Santra
Prof. Dr. Daniela Pfannkuche

Vorsitzende/r der Prüfungskommission: Prof. Dr. Robin Santra

Datum der Disputation: 21st November 2025

Datum der Druckfreigabe: 26th November 2025

Vorsitzender des Fach-Promotionsausschusses PHYSIK: Prof. Dr. Wolfgang J. Parak

Leiter des Fachbereichs PHYSIK: Prof. Dr. Markus Drescher

Dekan der Fakultät MIN: Prof. Dr.-Ing. Norbert Ritter

Abstract

The quest to understand the chemical processes that govern natural phenomena has been a central theme in scientific research throughout the history of humanity. From the large scales studied in astronomy and atmospheric sciences to the intrincate systems in biochemistry, all such processes ultimately trace back to the dynamics of molecules at atomic and subatomic levels. Quantum mechanics marked a significant breakthrough in our ability to describe and predict molecular dynamics compared to the previous empirical or statistical approaches that previously dominated the field of chemistry.

The central equations of molecular quantum mechanics at non-relativistic regimes are the Schrödinger equations. A range of theoretical and computational challenges arise when attempting to solve the Schrödinger equation. Among these challenges, the most prohibitive is the high-dimensionality inherent to the systems, which results from the necessity of characterizing the degrees of freedom for every constituent particle. This high-dimensionality leads to an extremely unfavorable scaling of the computational costs with the system size, a phenomenon commonly referred to as the curse of dimensionality. This thesis focuses on developing numerical methodologies that allow scalable computations. The content is divided in two different blocks. First, the time-independent vibrational Schrödinger equation will be adressed. Then, we explore the dynamics of molecules in the strong field physics limit. Concretely, we investigate the laser-induced electron diffraction imaging technique.

For the time-independent problem, we propose a methodology capable of creating parametrizable families of orthonormal basis sets. The creation of such sets is done by the introduction of a non-singular change of coordinates, modelled using normalizing flows (invertible neural networks). The normalizing-flows algorithm significantly enhances the approximation power of spectral methods by decoupling the different modes of the molecule. To obtain the optimal coordinate set, an optimization process of the approximated energies to the variational limit is performed. The use of optimization techniques— that only require the physical domain and an underlying basis set— significanlty reduces the required expertise for performing vibrational calculations. Based on an analysis of the optimal coordinates, we provide an explanation of how coordinate transformations are affected by basis set truncation, number of target states and coordinate domain. All this intuition is used to propose two different transferable capabilites of normalizing-flow coordinates: transferability across basis set truncations and across molecules with similar structural motif. The transferability across basis set truncations allows for an efficient training algorithm, and for the production of some of the most accurate vibrational results for highly-excited target states in the literature. Transferability across different molecules opens the possibility to create a collection of preoptimized coordinates that can be applied in a broad range of quantum-chemistry calculations.

In the second block, we aim to simulate the observables from the laser-induced electron diffraction imaging technique. In this case, the dynamics of the experiment are not solved by trying to find a direct solution of the time-dependent Schrödinger equation. Instead, we resolve the electron dynamics using the semiclassical model: an approximate model used in strong field physics that decomposes the continuum component of the wavefunction of an ionized molecule into individual classical photoelectron trajectories. The use of this approximation is motivated by the high-dimensionality of the systems, the complexity of the dynamics and long amplitude motions of photoelectrons. As part of the work, we implement all the necessary computational tools required to construct the semiclassical model: a tunneling ionization theory, approximations of the molecular electric field, electron trajectory propagators and a detection algorithm. As a single-active electron algorithm, the semiclassical model exhibits a favorable scaling with respect to the size of the molecule, making it suitable for the study of complex systems. We show that the algorithm is capable of replicating the experimentally measured momentum distributions, thereby validating its practical applicability.

Zusammenfassung

Eines der zentralen Themen in der wissenschaftlichen Forschung ist, wie chemische Prozesse Phänomene in der Natur bestimmen und gestalten. Die Dynamik von Molekülen auf atomarer und subatomarer Ebene bestimmt beispielsweise Eigenschaften auf großen Skalen, wie in der Astronomie und in den Atmosphärenwissenschaften, bis hin zu den komplexesten Interaktionen in der Biochemie. Eine quantenmechanische Beschreibung von Molekülen stellt einen bedeutenden Durchbruch in unserer Fähigkeit dar, die Dynamik von Molekülen zu beschreiben und vorherzusagen. Die grundlegenden Gleichungen in der Molekülphysik und der molekularen Quantenmechanik in nicht-relativistischen Domänen werden durch die Schrödingergleichung beschrieben. Das Lösen der Schrödingergleichung ist nicht trivial, da eine Vielzahl theoretischer und numerischer Herausforderungen überwunden werden muss. Unter diesen stellt die hohe Dimensionalität eines betrachteten Systems eine der größten Hürden dar, da die Freiheitsgrade jedes einzelnen Teilchens in einem System berücksichtigt werden müssen. Diese hohe Dimensionalität führt zu einer ungünstigen Skalierung der Rechenkosten mit wachsender Systemgröße – ein Phänomen, das allgemein als „Fluch der Dimensionalität“ bezeichnet wird. Das zentrale Thema dieser Dissertation befasst sich mit der Entwicklung neuer numerischer Methoden, die kostengünstige und skalierbare Berechnungen ermöglichen. Der Inhalt dieser Arbeit ist in zwei Teilgebiete unterteilt: Im ersten Teilgebiet wird die zeitunabhängige Schrödingergleichung im Rahmen molekularer Schwingungen betrachtet. Anschließend wird die Dynamik von Molekülen im Strong-Field-Regime untersucht, wobei der Schwerpunkt auf der laserinduzierten Elektronenbeugung liegt. Für das zeitunabhängige Problem wird im Rahmen dieser Arbeit eine Methodik vorgeschlagen, mit der sich parametrisierbare Familien orthonormaler Basissätze erzeugen lassen. Dies erfolgt durch die Einführung einer Koordinatentransformation, die mithilfe von Normalizing Flows (invertierbaren neuronalen Netzen) modelliert wird. Der Normalizing-Flows-Ansatz steigert die Näherungsgenauig spektraler Methoden erheblich, da er eine Entkopplung der verschiedenen Schwingungsmoden des Moleküls ermöglicht. Zur Bestimmung der optimalen Koordinaten wird ein Optimierungsverfah durchgeführt, das die approximierten Energien bis an das Variationslimit anpasst. Die Verwendung von Optimierungstechniken, die lediglich den physikalischen Definitionsbereich und einen zugrunde liegenden Basissatz erfordern, reduziert den notwendigen Aufwand für Rechnungen signifikant. Anhand der Analyse der optimalen Koordinaten wird erläutert, wie Koordinaten transformationen durch die Trunkierung des Basissatzes, die Anzahl der Zielzustände und den Koordinatenbereich beeinflusst werden. Diese Erkenntnisse führen zur Formulierung zweier übertragbarer Fähigkeiten in Normalizing-Flow-Koordinaten: die Übertragbarkeit über verschiedene Basissatztrunkierungen hinweg sowie zwischen Molekülen mit ähnlicher Struktur. Die erste ermöglicht effiziente Trainingsalgorithme und führt zu einigen der präzisesten in der Literatur berichteten Vibrationsresultate für hochangeregte Zielzustände. Die zweite eröffnet die Möglichkeit, eine Sammlung voroptimierter Koordinaten zu erstellen, die in einem breiten Spektrum quantenchemischer Berechnungen eingesetzt werden kann. Im zweiten Teil der Arbeit werden die Observablen der laserinduzierten Elektronenbeugungs-Bildgebung simuliert. Die Dynamik des Experiments wird hierbei nicht durch eine direkte Lösung der zeitabhängigen Schrödingergleichung beschrieben. Stattdessen erfolgt die Modellierung der Elektronendynamik mittels eines semiklassischen Ansatzes – eines in der Strong-Field-Physik gebräuchlichen Näherungsverfahrens, das die Kontinuumskomponente der Wellenfunktion eines ionisierten Moleküls in einzelne klassische Photoelektronenbahnen zerlegt. Die Verwendung dieser Approximation ist durch die hohe Dimensionalität der Systeme, die Komplexität der Dynamik sowie die großamplitudigen Bewegungen der Photoelektronen motiviert. Im Rahmen der Arbeit werden sämtliche notwendigen rechnerischen Werkzeuge zur Implementierung des semiklassischen Modells entwickelt: eine Theorie der Tunnelionisation, Näherungen des molekularen elektrischen Feldes, ein Algorithmus zur Propagation der Elektronenbahnen sowie ein Detektionsverfahren. Als Single-Active-Elektron-Algorithmus weist das semiklassische Modell eine günstige Skalierung in Bezug auf die Molekülgröße auf und eignet sich somit für die Untersuchung komplexer Systeme. Es wird gezeigt, dass der Algorithmus in der Lage ist, die experimentell gemessenen Impulsverteilungen zu reproduzieren, was seine praktische Anwendbarkeit belegt.

Declaration on 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.

Hamburg, 2025-11-26, (Álvaro Fernández Corral)

Contents

Acknowledgements

I would like to thank my supervisors, Prof. Jochen Küpper, Prof. Armin Iske and Dr. Andrey Yachmenev, for giving me the opportunity to continue my professional training. The past three years has been filled with learning and self-developement. Thank you for engaging in enlightening discussions about science, philosophy and lifestyle. From you, I have received new perspectives of the problems I encountered.

I would like to thank all the members of CMI for welcoming me, a theoritician, in a group full of experimentalists. I also want to thank the people that have left CMI but supported me at the beginning of my stay. From the doctoral students to the team leaders, I have felt embraced as part of the group. It is these people and the personal connections that make the collective effort of science make sense.

Most of my doctoral work has been performed in the frame of the CMI Theory team. I would like to thank the CMI Theory team members (a.k.a. the Robochimps): Andrey, Emil and Yahya. At the time of my arrival, normalizing flows was a new-born with a huge potential. Now, we have a pretty good understanding of the algorithm. The in-between was rough. We struggled to understand its functioning, implement it (we lost many fights to Jax), and to publish it. I struggled a lot. However, according to Andrey's philosophy, it is struggling that allows us to learn and, therefore, improve. In this sense, there is still a whole lot of struggling to be done together with the Robochimps. In the personal note, I would also like to thank the Robochimps for all the intelectual conversations, support and camaraderie. The time that you guys drove one hour to visit me in a hospital outside Hamburg few days before Christmas was very meaningful for me.

I would like to thank Nicolás Mendoza for his effort on the implementation and writing of the PINNs manuscript, and for the many conversations that have followed since. I am convinced that your future is going to be briliant.

Muchas gracias a Lucía por haber sido mi gran apoyo durante todo este tiempo. Gracias por acompañarme en esta aventura, escuchar mis quejas y ser mi zona de confort. Te he escuchado mil veces intentar explicar a tus familiares y alumnos en qué trabaja exactamente tu novio. Ahora que ya lo sabes explicar bien, es hora de buscar una nueva aventura. Espero tenerte de nuevo a mi lado en ella. Quiero agradecer a la familia de Lucía por su apoyo estos tres años. Intentaré no volver a llevármela tan lejos otra vez.

No he tenido la suerte de tener los otros pilares de mi vida, mi familia y amigos, tan cerca de mí. Quiero agradecer a mi madre por respaldarme y servirme continuamente de guía. Sobra decir que sin tu ayuda no hubiese llegado hasta aquí. A Manuel, gracias por tus visitas a Hamburgo. Espero que a ti y a Gema os hayan resultado cómodos mis sofás.

A todos los compañeros del baloncesto, a Miguel, Sergio y Rubén, muchas gracias por seguir recibiéndome con los brazos abiertos cada vez que vuelvo, como si el tiempo no hubiese pasado. Además de con los brazos, gracias Rubén por también recibirme con los melones abiertos.

Por último, y aunque dudo que lo llegue a leer, quiero agradecer a Mario Ballestero por ser el profesor que, a través de su carisma, me convenció de estudiar física. Me resulta difícil imaginar lo distinto que hubiese sido mi vida si hace 10 años hubiese tenido otro profesor.

Quiero dedicar esta tesis a mi padre, que no tuvo la oportunidad de ver cómo me desarrollaba hasta convertirme en físico, y a mi abuela Dolores, que lamentablemente se pierde la ocasión de ver a su segundo nieto doctor.

1 Fundamental concepts

1.1 The Schrödinger equation

For systems with sizes ranging from molecular to subatomic, quantum mechanics is the theory that provides the most accurate description. Unlike classical mechanics, where physical properties of objects are assigned definite values, quantum mechanics describes these properties using probability distributions. These distributions are encoded in the wavefunction, denoted by Ψ. The probabilistic knowledge of the observables of the system is totally contained within the wavefunction, allowing the reproduction of the statistical properties of the system from it. For a system consisting of N particles, the wavefunction is a complex-valued function defined in a $4N+1$-dimensional domain in position or momentum configuration space. In the position space, the wavefunction is

$$\Psi(\mathbf{r}_1, \dots, \mathbf{r}_N, \chi_i;\ t), \tag{1.1}$$

where $\mathbf{r}_i$ represents the three-dimensional position of the i-th particle, χ_i denotes its spin state, and t is the time. In this thesis, the spin properties of molecules will not be studied, and thus we will systematically omit this property throughout the text. We redefine the wavefunction as having a $3N + 1$-dimensional domain.

Wavefunctions are spatial square-integrable functions that satisfy the condition

$$\|\Psi\|_2^2\,(t) = \int_{\mathbb{R}^{3N}} d\mathbf{x}\ |\Psi(t, \mathbf{x})|^2 = 1, \tag{1.2}$$

where $\|\cdot\|_2$ denotes the L^2 norm. In this expression, the object $\Psi^*\Psi = |\Psi|^2$ mimics a probability density over space. This is the Born interpration of the wavefunction. The normalization condition$\|\Psi\|_2\,(t) = 1$ reflects that the probability of finding the particles in the domain is one at all times. This justifies the preferential use of normalized wavefunctions, i. e., wavefunctions of norm one.

The partial differential equation (PDE) that governs the evolution of the wavefunction in space and time is the time-dependent Schrödinger equation (TDSE) [1]

$$i\frac{\partial \Psi}{\partial t} = H\Psi, \tag{1.3}$$

where H is the Hamiltonian operator of the system. Throughout the thesis, we simplify the notation by adopting atomic units, i. e., $\hbar = e = m_e = 4\pi\epsilon_0 = 1$.

The Hamiltonian operator encapsulates the interactions between different particles, and between particles and the environment. The Hamiltonian is composed of the kinetic and potential operators [2, 3] and is typically written as

$$H = T + V, \tag{1.4}$$

where T is the kinetic energy operator and V is the potential energy operator.

The kinetic operator is a second-order differential operator that measures the kinetic energy of the particles. For a system of N particles, it is defined as

$$T = -\sum_{i=1}^{N} \frac{1}{2m_i} \frac{\partial^2}{\partial \mathbf{r}_i^2}, \tag{1.5}$$

where m_i is the mass of the i-th particle, $\mathbf{r}_i$ its position in Cartesian coordinates, and $\frac{\partial^2}{\partial \mathbf{r}_i^2}$ is the second-order derivative with respect to its position.

The potential operator is a multiplicative (functional) operator that measures the potential energy of the configuration of particles and the interaction with their environment at all times. Due to its multiplicative nature, the potential operator commutes with the wavefunction, while the kinetic operator generally does not.

If the potential operator does not depend explicitly on time, the Hamiltonian becomes explicitly time-independent, i. e., its time-dependency is implicit on the spatial configuration. Although the wavefunction still evolves in time, its dependency in time and space is separable in this case. Applying separation of variables to the TDSE, the time dependence can be factored out as a complex exponential. The wavefunction can be expressed as

$$\Psi(\mathbf{r}_1, \ldots, \mathbf{r}_N; t) = \sum_n c_n \psi_n(\mathbf{r}_1, \ldots, \mathbf{r}_N)\, e^{-iE_n t}, \tag{1.6}$$

where c_n are complex-valued constants determined by initial conditions and E_n are real constants that measure the frequency of complex-plane rotations. Each of the spatial components of the wavefunction, ψ_n, is stationary.

By introducing the ansatz (1.6) into the TDSE and isolating terms for each value of n, the time-independent Schrödinger equation (TISE) [1] is obtained

$$H\psi_n = E_n \psi_n. \tag{1.7}$$

The TISE is an eigenvalue equation on partial derivatives. It admits infinitely many pairs of eigenvalues and eigenfunctions (E_n, ψ_n). The constancy of the eigenvalues E_n arise as consequence of a symmetry in the Hamiltonian, specifically, invariance to time-translation. According to Noether's theorem [4], this symmetry implies the conservation of energy. Consequently, the eigenvalues of the TISE are denoted as energy states, or simply energies, and the eigenfunctions are the associated stationary states.

For any physically meaningful Hamiltonian, the energies are bounded from below [5]. The smallest energy of the system, E_0, is called the ground state energy and the eigenfunction associated to it, ψ_0, is the ground state. All other solutions are referred to as excited states and energies.

1.1.1 The Born-Oppenheimer approximation

In molecular physics, the systems are composed of electrons and nuclei. Following standard notation, the three-dimensional position of the i-th electron is denoted by $\mathbf{r}_i$ and the three-dimensional position of the α-th nucleus by $\mathbf{R}_\alpha$. The Hamiltonian operator can be separated into the kinetic operators of each set of particles and their interactions as [2]

$$H = T_e + T_N + V_{e-e} + V_{N-N} + V_{e-N}, \tag{1.8}$$

where e denotes electrons and N denotes nuclei. T_e and T_N are the kinetic energy operators of the sets of electrons and nuclei, and are defined as in Equation (1.5). V_{e-e} models the electron-electron interaction according to their relative position, V_{e-N} models the electron-nucleus interaction and V_{N-N} models the nucleus-nucleus interaction. For a system with N_e electrons and N_N nuclei, these potentials are given by

$$V_{e-e}(\mathbf{r}_i) = \sum_{1\leq i<j}^{N_e} \frac{1}{\left|\mathbf{r}_i - \mathbf{r}_j\right|}, \qquad V_{e-N}(\mathbf{r}_i, \mathbf{R}_\alpha) = \sum_{i=1}^{N_e} \sum_{\alpha=1}^{N_N} - \frac{Z_\alpha}{\left|\mathbf{r}_i - \mathbf{R}_\alpha\right|}, \tag{1.9}$$

$$V_{N-N}(\mathbf{R}_\alpha) = \sum_{1\leq\alpha<\beta}^{N_N} \frac{Z_\alpha\, Z_\beta}{\left|\mathbf{R}_\alpha - \mathbf{R}_\beta\right|},$$

where Z_α refers to the charge of the α-th nucleus.

Due to the large difference in the magnitude of their masses, $m_N > 10^3\, m_e$, the timescales of electronic displacements are orders of magnitude smaller than those of nuclei. This observation is core to the Born-Oppenheimer approximation [6], which allows one to treat the nuclear positions as fixed when solving the electronic part of the Schrödinger equations. The justification of this approximation is that the electronic response to changes in the nuclear configuration is almost instantaneous on the timescales of nuclear displacements. As a result, the molecular wavefunction can be approximately factorized into the nuclear and electronic parts as

$$\psi(\mathbf{r}_i, \mathbf{R}_\alpha) \approx \psi_e(\mathbf{r}_i; \mathbf{R}_\alpha) \cdot \psi_N(\mathbf{R}_\alpha; \psi_e). \tag{1.10}$$

Here, the electronic wavefunction depends parametrically on the position of the nuclei and the nuclear wavefunction depends implictly on the electronic configuration. Using this ansatz, both the TDSE and TISE can each be divided into two sets of equations. For the TISE, this separation leads to [2]

$$\left[T_e(\mathbf{r}_i) + V_{e-e}(\mathbf{r}_i) + V_{e-N}(\mathbf{r}_i; \mathbf{R}_\alpha) + V_{N-N}(\mathbf{R}_\alpha)\right] \psi_e^{(l)}(\mathbf{r}_i; \mathbf{R}_\alpha) = \hat{E}^{(l)}(\mathbf{R}_\alpha) \psi_e^{(l)}(\mathbf{r}_i; \mathbf{R}_\alpha), \tag{1.11}$$

$$\left[T_N(\mathbf{R}_\alpha) + \hat{E}^{(l)}(\mathbf{R}_\alpha; \psi_e^{(l)})\right] \psi_N^{(l,k)}(\mathbf{R}_\alpha) = E^{(l,k)} \psi_N^{(l,k)}(\mathbf{R}_\alpha; \psi_e^{(l)}), \tag{1.12}$$

where the superscript (l) denotes the l-th electronic eigenstate and (l, k) refers to the k-th nuclear eigenstate associated to the electronic state (l). Each of these superlabels can be substituted by the corresponding quantum numbers of the eigenstates. The dependence of each operator and wavefunction on the nuclear and electronic positions is written explicitly for easier comprehension.

Equation (1.11) models the wavefunction of electrons. Its dependence on the position of nuclei is treated parametrically, meaning that a solution can be found for each static choice of the nuclear positions $\mathbf{R}_\alpha$. Consequently, both the electronic wavefunction and the electronic energies depend parametrically on the choice of the nuclear positions that is introduced for the calculation.

Equation (1.12) models the wavefunction of the nuclei of the molecule. $\hat{E}^{(l)}(\mathbf{R}_\alpha; \psi_e^{(l)})$ represents the parametric energy obtained from solving the electronic problem and depends on the choice of electronic eigenstate $\psi_e^{(l)}$. The positions $\mathbf{R}_\alpha$ are no longer parameters but the coordinates of the nuclear problem. For this reason, a continous representation of $\hat{E}^{(l)}(\mathbf{R}_\alpha)$ is required. This is achieved by solving the electronic TISE at many different values of the $\mathbf{R}_\alpha$ and interpolating over them. This continuous representation of the electronic energies is usually refered to as the potential-energy surface (PES) [7, 8]. The interpolation and extrapolation quality of the PES is of great importance for reliable predictions of molecular energies and dynamics.

Many of the molecular dynamics of interest occur at relatively low temperatures. In this regime, the energy scale of the system is smaller than the gap between electronic states, making the electronic population mostly localised in the ground state. Even for systems at high temperatures, the electronic distribution can be mostly populated at the ground state. For this reason, disregarding the action of higher-lying electronic states is often a good approximation. The ground electronic state is sufficient for describing the systems in this thesis. From this point on, the PES is redefined as the electronic ground state PES,

$$V(\mathbf{R}_j) \equiv \hat{E}^{(0)}(\mathbf{R}_j). \tag{1.13}$$

For both electronic and nuclear TISEs, the classification of states depend on the relation between the energy and the values of the potential operator. A state is called a bound state if its energy is greater than the values of the potential operator in a *finite* subdomain [3, 9]. The probability density associated to a bound state is confined to a finite region. Otherwise, a state is referred as unbound or continuum state if its energy is bigger than the potential operator in an *infinite* subdomain. The density of continuum states is not confined to a finite region and they might be found arbitrarily far from the interaction center. Continuum states are typically associated with ionization (for electrons) or dissociation (for nuclei).

For bound states, there exists a discrete gap between every pair of energies. This makes the bound spectrum discrete and countable, or in other words, quantized. In contrast, the energy spectrum of continuum states is continuous, meaning that the eigenstates are dense in energy. By definition, the energies of bound states are smaller than the energies of all continuum states.

In electronic calculations, a fixed reference frame is intrinsically chosen by specifying the nuclear positions and the relative position of electrons with respect to them. However, in nuclear calculations, no reference frame is intrinsicly defined [10]. From a physical standpoint, molecular properties are invariant to overall translations of the center of mass (COM). For this reason, the three dimensions that correspond to displacements of the COM are usually neglected. Mathematically, the COM coordinates only appear on the kinetic energy term, making the PDE for these coordinates separable. Similarly, overall rotations of the inertial frame do not alter scalar observables like energies. However, the rotational degrees of freedom are always coupled to the internal coordinates in the operators. This coupling introduces rotational-vibrational interactions that can shift and split degenerate energy levels for rotating molecules.

For a a non-rotating molecule composed of N nuclei, $3N-6$ degrees of freedom are required [10]. These coordinates are refered to as the vibrational or internal coordinates of the nuclear motion and denoted by $\mathbf{x}$. The kinetic term of the Hamiltonian needs to be transformed according to the change of coordinates from Euclidian to internal coordinates. A general version of the nuclear TISE in any set of vibrational coordinates is

$$\left[\sum_{i,j} \frac{\partial}{\partial x_i} G^{ij}(\mathbf{x}) \frac{\partial}{\partial x_j} + U(\mathbf{x}) + \hat{E}^{(l)}(\mathbf{x})\right] \psi_N^{(l,k)}(\mathbf{x}) = E^{(l,k)}\, \psi_N^{(l,k)}(\mathbf{x}), \tag{1.14}$$

where $G^{ij}(\mathbf{x})$ is the mass-weighted metric tensor and $U(\mathbf{x})$ is the pseudopotential, arising from the cordinate transformation. These operators account for the transformation of the $3N$-dimensional Euclidian Laplacian operator to the $(3N-6)$-dimensional Laplacian operator in internal coordinates, ignoring the translational and rotational degrees of freedom. The details on the derivation and analytical expression of these operators can be found in Bunker and Jensen [10]. In addition, a symmetric (bilinear) formulation of the obtained kinetic operator can be achieved by using the Podolsky formulation of the vibrational Hamiltonian [11].

If rotational states of the molecule are explored, the 3 dimensions corresponding to the overall rotations need to be included and thus the Schrödinger equation is defined in a $(3N-3)$-dimensional domain. These coordinates are referred to as the ro-vibrational coordinates and are described by the set $(\mathbf{x}, \Gamma)$, where Γ are the 3 Euler angles that parametrise the three-dimensional rotations of molecules. The shape of the TISE in this coordinate set is the same as Equation (1.14), but with the introduction of the additional rotational coordinates as variables.

1.2 Mathematical insights of the Schrödinger equation

In the previous section, an introduction of the Schödinger equation was given without a detailed analysis of its mathematical properties. This introduction serves to understand the setup of the Schrödinger equations and the physical meaning of the wavefunction. However, to identify the challenges of finding its solution, a more rigorous mathematical analysis of the space of solutions is necessary.

1.2.1 The Hilbert space of the Schrödinger equation

For a system with N particles, the spatial component of wavefunctions ψ belongs to the manifold of square integrable functions defined in the spatial domain $\Omega \subseteq \mathbb{R}^{3N}$, that is $\psi \in L^2(\Omega)$. The definition of this space is given by

$$L^2(\Omega) := \left\{ f \mid f : \Omega \to \mathbb{C} \quad \text{and} \int_\Omega d\mathbf{r}\, |f|^2 < \infty \right\}, \tag{1.15}$$

where $\mathbf{r}$ denotes the variables of Ω. It can be shown that this set defines an infinite-dimensional Banach vector space [5] equipped with the L^2 norm. Additionally, this space admits an inner-product, denoted as $\langle \cdot | \cdot \rangle : L^2 \times L^2 \to \mathbb{C}$, and defined as

$$\langle f | g \rangle = \int_\Omega f^* g \; d\mathbf{r}, \tag{1.16}$$

where f^* is the complex conjugate of the function f. By its definition, the inner-product satisfies the following conditions

$$\begin{aligned} &\langle f|g\rangle = (\langle g|f\rangle)^* \quad \text{conjugate symmetry,} \\ &\left\|f\right\|_2^2 = \langle f|f\rangle = 0 \;\Leftrightarrow\; f = 0 \quad \text{positive definiteness,} \\ &\langle f|g\rangle \leq \left\|f\right\|_2 + \left\|g\right\|_2 \quad \text{Cauchy-Schwarz inequality.} \end{aligned} \tag{1.17}$$

Equipped with such inner-product, L^2 becomes a Hilbert space [5], denoted by $\mathcal{H}$. Following the inner-product definition, we adopt the bra-ket notation, where the ket $|g\rangle$ denotes a function in L^2 and the bra is a linear functional operator $\langle f| : L^2(\Omega) \to \mathbb{C}$ that acts on a ket $|g\rangle$ via the inner product $\langle f|g\rangle$.

An infinite set of functions $\left\{\phi_i\right\}_{i=0}^{\infty}$ is said to be a complete basis set of L^2 if any function $f \in L^2$ can be expressed as a linear combination of the functions in the basis, that is

$$f = \sum_{i=0}^{\infty} c_i \phi_i, \quad \text{for all } f \in L^2. \tag{1.18}$$

The set $\left\{\phi_i\right\}_{i=0}^{\infty}$ forms a Schauder basis set of L^2 if the coefficients c_i can be univocally determined.

If the inner-product of the functions in the basis set satisfies

$$\langle \phi_i | \phi_j \rangle = K_i \delta_{ij}, \tag{1.19}$$

where δ_{ij} is the Kronecker-delta, the basis is said to be orthogonal. Moreover, if all functions of the basis set are normalized, i. e., $K_i = 1$ for all i, then the basis is orthonormal. For orthonormal basis sets, the coefficients of the expansion can be found by simply projecting the function f onto the basis, i. e.,

$$c_i = \langle \phi_i | f \rangle. \tag{1.20}$$

Therefore, all orthonormal basis sets are Schauder basis sets.

Note that this way of finding the expansion coefficients by a projection requires prior knowledge of f, which in general will not be available.

By the expansion in Equation (1.18), the basis set can be used to compute the observables represented by operators. For any linear operator O, its action on $|f\rangle$ can be computed as a linear combination of the action of O on the basis set functions, that is

$$O|f\rangle = \sum_{i=0}^{\infty} c_i \, O|\phi_i\rangle, \tag{1.21}$$

where the linearity of O allows the term-wise action on the expansion. A basis set creates a useful representation for performing calculations of operators acting on the functions of the space and forms the foundation of most spectral and variational methods in quantum mechanics [2, 12].

The action (or recursive action) of operators onto $\mathcal{H}$ is not necessarily contained in itself. Moreover, the eigenfunctions of self-adjoint operators are not necesarily square-integrable functions. For this reason, the general setting for problems in quantum physics is rigged Hilbert spaces [13]. Rigged Hilbert spaces include two additional functional spaces: one extends the functional space of test functions to contain non-square-integrable functions. The second is the space of physically-aceptable functions, and is defined as the subset of $\mathcal{H}$ for which any combination of the action of operators is square-integrable. Both of these spaces are invariant under the action of operators, meaning that any operator acting on them yields a function that still belongs in the space. The definition of the rigged Hilbert spaces associated to a system is quite intrincate. In practice, numerical approximations of the wavefunction are assumed to belong to the associated physical subspace, rather than imposed.

1.2.2 Hamiltonian operator

The Hamiltonian operator H is a self-adjoint, unbounded operator defined on the domain $D(H) = H^2(\Omega) \subset L^2(\Omega)$, the Sobolev space of functions that together with their first and second order partial derivatives are square-integrable [5]. This domain is typically dense in L^2 [2, 12].

The self-adjointness requires that the Hamiltonian operator satisfies

$$\langle f|Hg\rangle = \langle Hf|g\rangle. \tag{1.22}$$

The unboundedness of the Hamiltonian operator means that its norm diverges over its domain

$$\|H\|_{D(H)} := \sup_{f\in D(H)} \frac{\|Hf\|}{\|f\|} \to \infty, \tag{1.23}$$

where $\|\cdot\|_{D(\cdot)}$ denotes the norm of an operator over its domain.

The unbounded nature of the Hamiltonian operator complicates the analysis of its mathematical properties. In particular, this means that there exist functions $|f\rangle \in D(H)$ for which $H|f\rangle$ does not belong in L^2.

Self-adjoint operators are characterized by having a real-valued spectrum, i. e., the eigenvalues of the operator are real. This explains why the observables in quantum physics are characterized by self-adjoint operators, as only real-valued measurements are achievable. In particular, this means that the spectrum of the Hamiltonian operator $\sigma(H) \subset \mathbb{R}$. The spectrum of the Hamiltonian is bounded from below by the ground-state energy and unbounded from above.

The set of eigenfunctions of the Hamiltonian defines an orthonormal basis set of the Hilbert space. The theorem that guarantees this property is the spectral projection theorem of self-adjoint operators in its projection-valued form. The proof of this theorem is intrincate. For this reason, we only present its enunciate. The reader is referred to Reed and Simon [12, Chapter VIII: Theorem VIII.6] for a more rigurous description of the theorem and its proof.

The spectral projection theorem states that an unbounded self-adjoint operator defined on a Hilbert space can be represented as the projection-valued integral over its spectrum. The Hamiltonian operator satisfies these conditions at all times. For this reason, the theorem is enunciated for a fixed time and the time-dependency is omitted. The projection-valued integral of the Hamiltonian spectrum is

$$H = \int_{\sigma(H)} E \, dP_E, \tag{1.24}$$

where P_E is the projection-valued measure that decomposes the Hilbert space into orthogonal subspaces associated with different spectral components of H.

In general, the spectrum of the Hamiltonian can be divided into the combination of the point spectrum $\sigma_p = \{E_i\}_{i=0}^{N-1}$, corresponding to discrete bound states, and the continuous spectrum $\sigma_c \subseteq [E_N, \infty)$, which is dense and corresponds to the continuum states. For a Hamiltonian whose spectra contains N bound states with eigenpairs $\{(\psi_i, E_i)\}_{i=0}^{N-1}$ and continuum states with energy E and eigenfunctions denoted by $\psi_c(E)$, the operator action can be written as the combination of the projections

$$H = \sum_{i=0}^{N-1} E_i|\psi_i\rangle\langle\psi_i| + \int dE \, E \, |\psi_c(E)\rangle\langle\psi_c(E)|. \tag{1.25}$$

This implies that the action of the Hamiltonian on any wavefunction can be decomposed into the sum of the actions of the projections of each eigenfunction. As previously discussed, this is in fact the fundamental property of basis sets. As a result of the spectral projection theorem, the eigenfunctions of the system form an orthonormal basis of the space $D(H)$, often dense in $L^2(\Omega)$. This will be a crucial feature in the construction of basis sets for numerical applications.

The spectral projection theorem does not guarantee the uniqueness of the eigenfunctions associated to an eigenvalue of the Hamiltonian. Eigenfunctions that share the same energy are called degenerate states and the corresponding energy is said to be a degenerate energy level.

The structure of the Hamiltonian spectrum depends on the physical system. Some systems may have purely discrete spectra—producing an infinite countable set—, while others may only exhibit continuous spectra. The quantum harmonic oscillator is a prototypical example of a system with purely discrete spectrum. Its eigenfunctions, i. e., the Hermite functions $\{h_n\}_{n=0}^{\infty}$, form a countable orthonormal eigenbasis. The energy associated to the n-th eigenfunction is $E_n = n + \frac{1}{2}$. The action of the Hamiltonian in the n-th basis function is $H|h_n\rangle = E_n|h_n\rangle$. The sequence $\{E_n|h_n\rangle\}_{n\to\infty}$ does not converge in L^2, highlighting that the image of the Hamiltonian of the harmonic oscillator is unbounded. This example illustrates that the creation of an orthonormal basis set from the eigenfunctions of a Hamiltonian is possible, even when its spectrum is not upper bounded.

1.3 Spectral methods

Equation (1.21) shows that the action of any linear operator can be represented by the projection onto the functions in a basis set. Using the representation of the Hamiltonian in a basis set, solutions of the Schrödinger equations can be approximated. An algorithm that uses the functions of a basis set to approximate functions of the same functional space is called a spectral method. In this section, we explore the application of spectral methods to the TISE and TDSE.

1.3.1 Spectral methods for the TISE

Given a Hamiltonian operator H that defines a TISE in the domain $\Omega \subseteq \mathbb{R}^D$, its action on the functions of a basis set $\left\{\phi_j\right\}_{j=0}^{\infty}$ of L^2 is $H|\phi_j\rangle$. Taking the inner-product of the result with another basis function ϕ_i yields

$$H_{ij} := \langle\phi_i|H\phi_j\rangle, \tag{1.26}$$

where $\mathbf{H} = H_{ij}$ is the matrix representation of the Hamiltonian in the span of the basis set functions. $\mathbf{H}$ is a complex-valued unbounded Hermitian matrix, $H_{ij} = H_{ji}^*$.

Let ψ_k be any normalized eigenfunction of H with eigenvalue E_k. This function satisfies the condition $\langle\psi_k|H\psi_k\rangle = E_k$. This integral formulation of the Schrödinger equation is called the weak formulation of the TISE, as it seeks a function whose energy expectation value is equal to its eigenvalue. This is a weaker condition than the original pointwise equation, also called the strong formulation of the TISE.

If the chosen basis set is complete, the eigenfunctions of the Hamiltonian are contained in the closure of the span of the basis. Solving the matrix eigenvalue problem of $\mathbf{H}$ yields the exact eigenvalues of the TISE. Specifically, if the eigenfunctions are resolved in the basis set by $\psi_i = \sum_{j=0}^{\infty} c_{ij}\phi_j$, the diagonalizion of the Hamiltonian matrix produces

$$H_{ij} = \sum_{k,l=0}^{\infty} c_{ik}^{*}\, E_{kl}\, c_{lj}, \tag{1.27}$$

where $\mathbf{E}$ is a diagonal matrix with the energies of the system as its entries, $E_{kl} = E_k \delta_{kl}$ and $\mathbf{c}$ is a unitary matrix containing the coefficients of the expansion. This construction ensures that all states $\{\psi_i\}_{i=0}^{\infty}$ are orthonormal, provided $\{\phi_i\}_{i=0}^{\infty}$ is orthonormal.

Due to the infinite vector space nature of the Hilbert space, any numerical representation must be truncated and thus cannot be complete. In practice, spectral methods use N basis set functions to approximate eigenfunctions and eigenvalues. This leads to approximated energies $\hat{E}_i$ obtained from the diagonalization of the Hamiltonian matrix in the span of the N functions, $\hat{\mathbf{H}} = \mathbf{c}^{\dagger}\ \hat{\mathbf{E}}\ \mathbf{c}$, where $\mathbf{c}$ is a unitary matrix. Using this construction, the corresponding approximated eigenfunctions are given by

$$\hat{\psi}_i := \sum_{j=0}^{N-1} c_{ij}\phi_j. \tag{1.28}$$

In most applications, the goal of the spectral method is to approximate a subset of M solutions of the TISE, and not infinitely many of them. In this case, not using a complete representation of the functional space does not prevent high accuracy of the calculations. When using N functions of a basis set ($N \geq M$), exact solutions are achieved when the eigenfunctions are contained in the linear span of the functions in the truncated set

$$\{\psi_i\}_{i=0}^{M-1} \subset \operatorname{span}\left(\{\phi_i\}_{i=0}^{N-1}\right). \tag{1.29}$$

This condition determines the efficiency of the spectral method. If it can be achieved with $N \approxeq M$, the basis is well-adapted for the problem and the spectral method is optimal. The similarity between the basis functions—or more accurately, of their linear span— and the eigenfunctions of the Hamiltonian is a desired condition. The choice of a basis set plays a central role in determining the efficiency of the spectral method.

This rationale also highlights the adequacy of spectral methods for solving the TISE. As the solution of the TISE is not unique, a spectral method can approximate many solutions at the same time. When using N basis functions, the numerical expense of building the Hamiltonian matrix is independent of the number of approximated eigenfunctions M. The overall computational cost can only be reduced by decreasing M using advanced diagonalization techniques. For this reason, more than one target state is very often chosen for the approximation.

1.3.2 The variational limit

The variational limit builds upon the spectral projection theorem and provides an understanding of the convergence of approximated energies to the eigenvalues of the Hamiltonian. In particular, it explains the convergence of a spectral method to the solutions of the TISE.

Consider the application of any normalized function $|f\rangle$ onto a Hamiltonian, which results in some approximated energy, $\hat{E} = \langle f|Hf\rangle$. By the spectral projection theorem, the set of eigenfunctions of the Hamiltonian defines an orthonormal basis set of L^2. As a matter of fact, f can be expanded as a linear combination of the eigenstates, which is equivalent to acting with $|f\rangle$ from the right in Equation (1.25). For simplicity of the proof, we will assume only bound states are solutions of the TISE and define the expansion in the eigenfunctions of H as $f = \sum_{i=0}^{\infty} c_i \psi_i$. Using this expansion, the approximated energy yields

$$\hat{E} = \sum_{i,j=0}^{\infty} \langle c_i \psi_i | H \, c_j \psi_j \rangle = \sum_{i=0}^{\infty} |c_i|^2 E_i \geq E_0, \tag{1.30}$$

where the orthonormality of the eigenfunctions was used for the second equality. The inequality uses the normalisation of the wavefunction, i. e. $\sum_{i=0}^{\infty} |c_i|^2 = 1$, and the fact that eigenvalues are lower-bounded by the ground state energy. Then, for any trial function, the obtained approximated energy can only be bigger than or equal to the ground state energy. The equality is achieved if and only if the trial function is exactly the ground state. This result is broadly known as the variational limit or variational principle [3, 9].

A variational limit for the first excited state E_1 can only be found by restricting the domain of trial functions to a subspace that is orthogonal to ψ_0, as in this space the energetic lower-bound is now E_1. By extension, a variational limit for the N-th excited state energy can be created by restricting the trial function to belong to the subspace that is orthogonal to all the $N-1$ lower-lying eigenfunctions.

Combining the set of variational limits for many functions, we arrive to the generalised variational limit. For any set of M orthonormal functions, $\left\{\phi_i\right\}_{i=0}^{M-1}$, the sum of the M approximated energies are bounded from below by the sum of the first M energies, i. e.,

$$\sum_{i=0}^{M-1} \langle \phi_i | H \phi_i \rangle = \sum_{i=0}^{M-1} \hat{E}_i \geq \sum_{i=0}^{M-1} E_i. \tag{1.31}$$

The equality in energies is again obtained only if the trial functions are the first M eigenstates, or a unitary linear combination of them (as the trace is invariant under such transformations). This implies that the approximated energies from the spectral method approach the system's energies from above as the number of basis set function increases.

Having a lower-bound for the approximated energies is the keystone of most of the variational optimisation processes for finding solutions of the TISE. Let $f(\theta)$ be a parametrised trial function , where θ denotes the set of all parameters. The best possible approximation to the ground state is obtained by minimising the approximated energy. In other words, to approximate the ground state, an optimisation can be performed to find the set of parameters θ_0 that minimises the approximated energy. Using the notation of machine learning algorithms, a loss function (the target function to be minimized) can be defined by direct application of the variational limit as

$$\mathcal{L}(\theta) = \langle f(\theta) | H f(\theta) \rangle. \tag{1.32}$$

Although $f(\theta_0)$ is the best possible approximation to the ground state within the parametrization space, it might or might not be a good approximation of the ground state. This depends on the capability of the ansatz to mimic the ground state.

This loss function can be extended to many trial functions by taking the sum of the approximated energies, that is, using the generalised variational limit. For all of our numerical approximations of the solutions of the TISE, the generalised variational limit (1.31) will be used as the loss function for the optimisation.

1.3.3 Spectral methods for the TDSE

In the previous two sections, we explored the utility of spectral methods applied to the TISE. The use of basis sets is not restricted to time-independent problems; it can also be benefitial in time-dependent dynamics. Since, at each time, the wavefunction is a square-integrable function, one natural strategy to solve the TDSE is to express the evolving wavefunction as a time-dependent expansion in a static spatial basis set.

Consider the TDSE governed by a time-dependent Hamiltonian $H(t)$ defined in the Hilbert space of square-integrable functions on the domain $\Omega \subseteq \mathbb{R}^D$. Let $\left\{\phi_i\right\}_{i=0}^{\infty}$ be a basis set of said functional space. The initial condition of the propagation is given by the projection onto the basis as

$$\Psi(\mathbf{r}, t=0) = \sum_{i=0}^{\infty} c_i^{(0)} \phi_i(\mathbf{r}), \tag{1.33}$$

where $c_i^{(0)}$ are the initial complex-valued coefficients.

We are interested in studying the dynamics of the wavefunction from this initial condition. Given that the basis set is a good representation of the spatial component, a common decision is to encapsulate the time-dependency of the wavefunction in the linear coefficients [2, 14]. This produces the ansatz

$$\Psi(\mathbf{r}, t) = \sum_{i=0}^{\infty} c_i(t) \phi_i(\mathbf{r}), \tag{1.34}$$

where $c_i(t=0) = c_i^{(0)}$.

Introducing this ansatz in the TDSE, we arrive to

$$i \sum_{i=0}^{\infty} \frac{\partial c_i(t)}{\partial t} \phi_i(\mathbf{r}) = \sum_{i=0}^{\infty} c_i(t) H(t) \phi_i(\mathbf{r}). \tag{1.35}$$

In this step, we have transformed a PDE that only required solving one function into a PDE that has up to infinitely-many functions to be determined. While this may seem like a disadvantage, we will show how this simplifies the calculations. By taking the inner product with the basis function $\langle \phi_j |$, we obtain

$$i \frac{\partial c_j(t)}{\partial t} = \sum_{i=0}^{\infty} c_i(t) \langle \phi_j | H(t) \phi_i \rangle, \tag{1.36}$$

where the orthonormality of the basis has been used in the left-hand side of the equation. In this step, the problem has been divided into a set of coupled PDEs. Importantly, the spatial variables of the problem are no longer contained in this equation, as they are integrated out.

In many molecular physics applications, the system starts from its equilibrium state and is perturbed by a time-dependent external field. This means that the Hamiltonian can be split into the molecular Hamiltonian H_0, which is time independent, and a time-dependent perturbation $h(t)$. The matrix obtained from the integral of the basis set functions with H_0 is the Hamiltonian matrix representation, previously defined as $\mathbf{H}$ in Equation (1.26). In the same spirit, we define the matrix obtained from the spatial integral of the perturbation Hamiltonian in a basis representation as $h_{ji}(t) = \langle \phi_j | h(t) \phi_i \rangle$.

The solution of Equation (1.36) is found by integrating the right-hand side and choosing the correct initial conditions, leading to

$$c_j(t) = \sum_{i=0}^{\infty} c_i^{(0)} e^{-i \int_0^t H_{ji} + h_{ji}(t')\, dt'}. \tag{1.37}$$

Analytically solving this equation guarantees an exact solution of the TDSE due to the completeness of the basis set. Once more, this is not feasible in numerical calculations and a set of N basis functions needs to be chosen for the numerical approximation. In that case, the spectral method yields an exact solution if and only if the dynamics are completely contained in the span of the N functions, even if the integration is performed analytically.

A numerical solution of (1.37) can be found by dividing the time interval t into sufficiently small steps Δt and approximating the integral between time steps using the trapezoidal rule, or other low-order integral approximations. This small time stepping is also benefitial for computing the exponential of the matrix, as only power series approaches provide good accuracy and, if the matrix entries are close to zero, only a low-order power series expansion is required.

The derivation of the solution using a basis set was performed for an arbitrary choice of basis set. A particularly convenient choice is to use the N first eigenfunctions of the TISE $\{\psi_i\}_{i=0}^{N}$, which satisfy the equation $H_0|\psi_i\rangle = E_i|\psi_i\rangle$. In this case, the Hamiltonian matrix $\mathbf{H}$ becomes diagonal $H_{ij} = E_i \delta_{ij}$. If no perturbation is present, the system will remain in the same state, varying only its phase at the rate of the energy. This is consistent with the separation of variables that was used for deriving the TISE from the TDSE.

When a time-dependent perturbation is present, it is the perturbation Hamiltonian $\mathbf{h}(t)$ that mixes the different eigenstates throughout the time evolution. In simulations of the dynamics under the action of an external field, the perturbation Hamiltonian is proportional to the dipole matrix elements

$$\mu_{ij,l} = \langle \psi_i | x_l \psi_j \rangle, \tag{1.38}$$

where x_l are the Cartesian coordinates in the laboratory frame, $l = x, y, z$. Due to its direct role in driving transitions between states, the dipole matrix is a fundamental observable in molecular dynamics and spectroscopy.

The choice of the eigenfunctions of the static Hamiltonian as a basis for the problem is sensible as long as the entries of H_{ij} dominate over those of the perturbation h_{ij}. This is the weak field regime. In contrast, if the action of the perturbation is comparable or stronger than the static Hamiltonian, the evolution may not be well represented by the stationary states. This is the strong field regime and is typically found in ionisation/dissociation processes. An alternative choice of basis set in the strong field regime is the set of eigenfunctions of the perturbation Hamiltonian $h(t)$. With this choice, it is the molecular Hamiltonian H_0 that serves as a perturbation that mixes different functions.

This highlights one common misconception: the need for explaining the dynamics as projected onto the stationary states. During the dynamics of strong field processes, the population of bound states evolves very quickly. As a consequence, bound states do not provide a meaningful description of the system at each time. Moreover, continuous local deformations of the spatial component of the wavefunction may not be well-captured in the linear span of bound states. Accurately describing these transformations might require a large number of eigenstates. This problem can be identically found on other choices of static basis sets. However, only the molecular eigenbasis is used to *explain* the dynamics of the molecule in time.

Spectral methods for the TDSE with static basis sets face significant challenges in representing continuum states. This problem can be illustrated in the example of simulating dissociation dynamics. To accurately capture dissociation dynamics, all bound states with significant population at any time need to be included in the calculation. If there is no prior knowledge of the dynamics that allows us to discard certain states, all bound states might need to be included in the propagation. This creates a bottleneck on both solving the TISE to resolve many highly-excited states, and in the calculation of the perturbation Hamiltonian, which must be computed at all time steps for a great number of functions. Moreover, to allow for dissociation, continuum and resonant states (states with an imaginary energy component that decay into the continuum) need to be part of the calculations. As continuum states are dense in energy, including infinitely-many continuum states does not guarantee the convergence of the method, as the particular states that model the dissociation channels might still be missing.

The combination of these problems with the great memory costs of calculations that use a large number of basis set functions are a major bottleneck for applying spectral methods to solving the TDSE.

1.3.4 Alternatives to spectral methods

For the reasons outlined above (and others beyond the scope of this dicsussion), alternative formulations using continuous or semi-continuous representations of the wavefunction have been explored in quantum molecular dynamics.

One possibility is to perform the simulation using parametrised spatial functions, that are not necesarilly orthogonal, allowing the parameters to evolve in time. The time evolution of the parameters is computed together with that of the coefficients of the linear expansion to solve the TDSE as accurately as possible. This is the approach of the multiconfigurational time-dependent Hartree method (MCTDH) [15, 16]. We denote the MCTDH approach as having a semi-continuous representation. This refers to the ability of the representation to evolve in time, creating a continuous spatial representation in the span of the parameter set, while having a discrete representation in time.

Two major problems are encountered when solving the TDSE using MCTDH. First, the flexibility of the approximating functions is limited due to the use of a basic parametrization. This constrains the expresivity of the method. On the other hand, in implementations such as the Gaussian MCTDH, no restriction is made on the parameters such that the functions retain orthonormality. This means that the functional set needs to be normalised at all time steps, requiring the inversion of the overlap matrix. Depending on the initial conditions, the evolution of the parameters might create nearly linearly-dependent functions. In this case, the determinant of the overlap matrix approaches zero, making its inversion numerically unstable. This makes MCTDH unstable and dependent on the choice of initial values of the parameters. Moreover, these instabilities on the overlap matrix make the propagation difficult to integrate [17].

A similar scheme to MCTDH that uses a time-dependent orthonormal basis set will be layed out in the subsection 7.2.2. This approach aims to alleviate these two major drawbacks.

The recent development of machine learning techniques has also provided new promising tools for solving the TDSE. Among these numerical approaches, physics-informed neural networks (PINNs) stand out as a fully continuous approach. In this method, time and space are treated identically as variables of the wavefunction, which is modelled directly by a neural-network.

The foundational idea of PINNs was proposed by Raissi et al. [18] and relies on the fact that neural networks are universal approximators of functions, i. e., any function can be approximated to any precision by a neural network. The TDSE, as any other general PDE, can be written in the generic form

$$\mathcal{N}(u, \partial_{\mathrm{x}} u, \partial_{\mathrm{x}}^2 u, \dots) = 0, \tag{1.39}$$

where u is the target function and $\mathbf{x}$ the variables of the problem, which include time and spatial variables. We use this generic representation of the TDSE to remark that the method is extendable to any choice of PDE. To map the TDSE to this form, substitute

$$\begin{aligned} u &= \Psi \\ \mathcal{N} &= i\frac{\partial \Psi}{\partial t} - H\Psi. \end{aligned}$$

A solution is approximated by $u \approx \hat{u}(\mathbf{x}; \theta)$, where $\hat{u}(\mathbf{x}; \theta)$ is a neural network that depends on the parameters θ. The parameters of the neural network are optimised to minimise the loss function

$$\mathcal{L}(\theta) = \left\| \mathcal{N}(\hat{u}, \partial_{\mathrm{x}} \hat{u}, \partial_{\mathrm{x}}^2 \hat{u}, \dots) \right\|. \tag{1.40}$$

Because the physical information of the problem is directly passed in the loss function, these constructions receive the name ”physics-informed”.

The PINNs approach is, in principle, applicable to both the TISE and TDSE. However, in the case of the TISE, the existence of infinitely many orthonormal solutions and unknown eigenvalues makes the application not easily compatible. In contrast, PINNs offer the potential for high-accuracy solutions for equations with a unique solution, such as the TDSE. While the set up looks ideal, PINNs have found severe difficulties in finding solutions to the TDSE in general systems due to the excesive flexibility and the complexity of the parameter landscape. The optimisation of the PINNs often gets stuck in local minima or just finds the naive solution $\Psi = 0$.

An example of the use of PINNs for solving PDEs in a very different setup can be found in Chapter 6. There, a general time propagator of explicitly time-independent systems is sought, which for a fixed time-step size is unique. This provides a good example of the application of PINNs for approximating a unique solution.

Last, in many molecular physics applications, the primary goal of the computational models is to interpret or predict the outcome of some experiment. Modern experimental setups often involve highly complex instrumentation and intricate molecular dynamics, which are difficult to replicate in an analytical representation. For this reason, an accurate description of the wavefunction at all times, i. e., an exact solution of the TDSE, is not accesible by the means of traditional calculations and, often, not even required. Rather, an estimation of the observables that are obtained as a result of the experiment is enough to provide the necessary information of the dynamics. Such results can be obtained by models that reduce the complexity of the system by applying several layers of approximations that reduce the available information, but that make the simulation possible. These models might seem oversimplified and a coarse approximation of the physics taking place, but often provide the only tool for interpretability and prediction.

An example of such an experiment-oriented simulation is the semiclassical model of the strong-field dynamics of laser-induced electron diffraction, described in section 5.2.

1.4 Constructing basis sets of L^2

Approximating the wavefunction in the span of a basis set is an effective strategy for solving the Schrödinger equations. This approach offers several advantages. The convergence of the approximation with respect to the number of functions is guaranteed. For the TISE, several solutions can be approximated at once. In this section, we outline the construction of basis sets for multidimensional problems, highlighting practical considerations and common strategies.

1.4.1 Direct tensor-product basis

Basis sets for multidimensional spaces are usually constructed from functions defined in reduced-dimensional Hilbert spaces. When the domain of the problem allows for a separable structure, a direct tensor-product basis is possible, simplifying the mathematical formulation and computational implementation of the approach.

Let $L^2(\Omega)$ be the Hilbert space of square-integrable functions on the domain Ω. If the domain is expressable as the Cartesian product of two subdomains $\Omega = \Omega_1 \times \Omega_2$, the corresponding Hilbert space admits a decomposition into a tensor product [12, 19]

$$L^2(\Omega) = L^2(\Omega_1) \otimes L^2(\Omega_2). \tag{1.41}$$

Under this decomposition, the inner-product between two tensor-product functions is separable

$$\langle x_1 \otimes y_1 | x_2 \otimes y_2 \rangle_{L^2(\Omega)} = \langle x_1 | x_2 \rangle_{L^2(\Omega_1)} \langle y_1 | y_2 \rangle_{L^2(\Omega_2)}, \tag{1.42}$$

where $x_1, x_2 \in L^2(\Omega_1)$ and $y_1, y_2 \in L^2(\Omega_2)$. This property naturally generalizes to higher-dimensional tensor product spaces.

In molecular physics, the domain of the problems $\Omega \subseteq \mathbb{R}^D$ is typically decomposable as a direct Euclidian product of one-dimensional independent subdomains

$$\Omega = \Omega_1 \times \cdots \times \Omega_D, \tag{1.43}$$

where $\Omega_i \subseteq \mathbb{R}$ for all i. This implies that each spatial coordinate spans a finite, semi-infinite or infinite segment of $\mathbb{R}$. For instance, the domain associated to the three-dimensional spherical coordinates is $[0, \infty) \times [0, \pi] \times [0, 2\pi) \subset \mathbb{R}^3$. Although this space physically represents a ball that covers $\mathbb{R}^3$ completely, the coordinate space is a rectangular cuboid (a three-dimensional parallelepiped that has rectangles as its faces). This illustrates that the physical interpretation and the representation of the coordinate domain are conceptually distinct.

Given a Hilbert space with such a separable structure, an orthonormal complete basis set can be created as a direct product of basis set functions of each subdomain [12, 19]. That is, a basis set function in D dimensions can be expressed as

$$\phi_i(\mathbf{r}) = \phi^{(1)}_{i_1}(r_1) \cdot \phi^{(2)}_{i_2}(r_2) \ldots \phi^{(D)}_{i_D}(r_D) \tag{1.44}$$

where $i = (i_1, i_2, \ldots, i_D)$ is a contracted index. $\phi^{(j)}_{i_j}(r_j)$ denotes i_j-th element of the orthonormal complete basis set of $L^2(\Omega_j)$. Each coordinate may be assigned a different basis set suited to the geometry or boundary conditions of Ω_j. The superindex (j) in the basis set functions refers to this choice.

The set $\{\phi_i\}_{i=0}^{\infty}$ yields a complete orthonormal basis set of $L^2(\Omega)$. This type of construction of basis sets is called a direct-product basis set of univariate functions and is the most used type of basis sets due to their simplicity.

A significant drawback of direct-product basis sets is their exponential scaling with the dimensionality. If n_i functions are used in each coordinate, the total number of multidimensional functions is $\prod_{i=1}^{D} n_i$. In the case that all dimensions share the number of functions $n_i = n$, the total number of functions is n^D, which quickly becomes intractable as D increases, a phenomenon known as the curse of dimensionality.

1.4.2 Univariate basis functions

To define multimensional basis sets using the product in Equation (1.44), the possibility of creating a range of univariate basis sets is needed. By the spectral projection theorem, the eigenset of a self-adjoint operator forms a complete basis set of the space. For this reason, one option for constructing basis sets is solving one-dimensional TISEs.

Many standard univariate basis sets originate from this approach. The Hermite functions are the solution of the quantum harmonic oscillator and the generalized Laguerre functions are a solution of the radial part of the electronic wavefunction of the hydrogen atom. Other basis functions such as the Legendre polynomials are the solution of some differential equation that can be casted into a Sturm-Liouville problem, resembling of the TISE. All of these examples share important structural properties, as they are formed as products of polynomials and a weight function. In addition, they satisfy a recurrence relation between succesive functions and they can be generated via Rodrigues' formula [20]. Last, they admit exact integration by Gaussian quadratures, as explained in section 1.5.1. A basis set that satisfies these properties is called a classical orthogonal polynomial basis set. Although other non-polynomial-based basis sets can be created as solutions of arbitrary TISEs, the lack of these useful properties make them generally impractical. For this reason, these classical orthonormal polynomial basis sets are usually preferred as underlying basis sets for calculations.

Another practical concern is the mapping between the domains of the one-dimensional basis Ω_ϕ and the corresponding coordinate domain Ω_i. In many cases, a linear transformation of the domains provide an acceptable mapping. When infinite domains are morphed into finite ones by linear maps, it is typically assumed that the basis functions decay before arriving to the boundary of the physical domain. This introduces a systematic approximation error that may affect accuracy or the capability of extending the calculations to many states. This can be mitigated by the introduction of non-linear domain mappings, which is a central topic of this thesis. An in-depth investigation on the application of non-linear mappings of domains is provided in the chapters 2 and 3.

1.4.3 Contraction scheme

The basis set structure guarantees the convergence of the spectral method. However, the exponential scaling with the problem dimensionality and the infinite nature of Hilbert spaces make accurate calculations not always numerically achievable in the span of a direct product basis set. To adress this, it is desirable to construct basis sets in close relation with the eigenfunctions of the system. This is the main idea behind the method called contraction. Contraction was firslty popularized by Carter and Handy [21]. Since then, contraction has been established as a standard technique for multidimensional vibrational calculations [22–24]. The contraction algorithm is described below.

Let a TISE be defined on a D-dimensional decomposable domain $\Omega = \Omega_1 \times \cdots \times \Omega_D$ with $\Omega_i \subseteq \mathbb{R}$. Let $\mathbf{r}$ denote the variables of the system, T be the kinetic operator and V be the PES. We omit the pseudopotential term, as its action can be added to the PES. Let the PES have its global minimum at the point $\mathbf{r}^0 = (r^0_0, r^0_1, \ldots r^0_D)$, known as the equilibrium position. For most molecular systems, the assumption of having a single global minimum is physically justified.

Then, a basis set for the i-th variable can be created by solving the univariate TISE defined by *freezing* all other coordinates to their equilibrium value

$$\left(T(r_i) + V(r^0_0, \ldots, r_i, \ldots r^0_D)\right) \psi^{(i)}_n(r_i) = E_n \psi^{(i)}_n(r_i), \tag{1.45}$$

where $T(r_i)$ refers to the kinetic operator acting on the i-th variable uniquely. The second-order derivatives involving r_i and a different variable are not included in $T(r_i)$.

Each of these univariate TISE is typically solved by its representation on a different univariate basis set—usually a classical orthogonal polynomial basis set—truncated at n_i functions. As a result of solving this one-dimensional TISE, m_i eigenfunctions are obtained, $m_i < n_i$, hence the name of the method is contraction. These newly calculated functions belong to a basis set that was built in close resemblance to the problem in this dimension. A multidimensional basis set can be created as a direct product of the collection of these contracted sets. This basis set provides a reduction to the number of total funcions that achieve comparable accuracy to the underlying basis set. The contraction step adds complexity to the construction of the multidimensional basis set. Nonetheless, the benefits of contraction outweight the additional effort.

Alternatively, the one-dimensional basis functions can be obtained from solving the TISE in the mean potential field of the ground state of the other coordinates, that is $V(r_i) = \langle \prod_{j\neq i} \psi^{(j)}_0(r_j) | V(\mathbf{r}) | \prod_{j\neq i} \psi^{(j)}_0(r_j) \rangle$. This is the approach used in vibrational self-consistent-field calculations (VSCF) [25–27]. VSCF iteratively updates the approximated ground state in each dimension until a stationary multidimensional solution is found. While the algorithms differ, the objective of contraction and VSCF is shared: obtaining an optimal basis set per dimension.

Within the contraction framework, other divisions of the multidimensional domain are possible. Concretely, the division can be performed into multidimensional subdomains $\Omega_i \subseteq \mathbb{R}^{d_i}$, where $\sum_i d_i = D$. Such divisions are specially useful for high-dimensional systems, where a direct product basis of univariate functions becomes unfeasible due to the exponential scaling, even where very few functions per dimensions are used.

Because multidimensional basis sets are rare (with spherical harmonics being a notable exception), the basis set of each subdomain is typically built using the direct product of d_i univariate basis sets followed by a contraction. The TISE is solved for the variables of the subdomain while *freezing* the other coordinates to their equilibrium position. For each subdomain, this algorithm produces m_i approximate eigenfunctions, substantially fewer than the original set of $\sim n^{d_i}$ functions per subdomain. The multivariate contraction drastically reduces the total number of functions by reducing the degree of the set.

At each contraction, the amount of information that is available from the original underlying basis set is reduced. However, the contraction is designed to retain the key linear combinations that contribute most to the system's eigenfunctions, while discarding the negligible components. By performing repeated low-dimensional diagonalisations and selectively truncating the eigenstates, the computational bottleneck of directly building and diagonalizing the high-dimensional Hamiltonian matrix is mitigated. The amount of information discarded on the low-dimensional diagonalization and truncation of eigenstates is comparable to that of truncation of states after diagonalization of direct-product sets, thereby justifying the preferential use of contraction methods.

The strategy used for partitioning the coordinate space in subdomains is crucial for the quality of the calculations. This choice is typically made based on the relation between coordinates in the Hamiltonian operator. Coordinates whose role in the Hamiltonian is separable or nearly separable from the rest of coordinates are usually split into their own unidimensional contraction, as their contribution to the eigenfunctions is well mimicked by a univariate basis set. Coordinates that are deeply related through the operators of the Hamiltonian, i. e., strongly coupled, are grouped together for a multivariate contraction. If strongly coupled coordinates are separated in different partitions, the contraction might not lead to improved convergence to the solution (compared to the underlying basis set).

1.5 Integration schemes

Throughout this chapter, we have consistently relied on the Hilbert space inner-product and orthonormality relations to define quantities in the representation of a basis. In many calculations, such as the evaluation of the Hamiltonian matrix H_{ij}, an analytical expression of the integral values is not available and we need to explore numerical solutions. In this section, we describe two of the most commonly used approaches: Gaussian quadratures and Monte Carlo integration.

1.5.1 Gaussian quadratures

Define the functions of an orthonormal basis set $\left\{\phi_i\right\}_{i=0}^{\infty}$ of L^2 in a unidimensional domain $\Omega \subseteq \mathbb{R}$ as the multiplication of a classical orthogonal polynomial p_i and the square root of a weight function w

$$\phi_i = p_i \sqrt{w}. \tag{1.46}$$

The weight function is positive ($w(x) > 0$ for all $x \in \Omega$), common to all functions in the basis and ensures the adequate decay at the domain boundaries such that the polynomials are square-integrable. In addition, this weight function is selected in close relation with the set of polynomials such that the Gaussian quadrature relation can be employed for accurate integration.

Gaussian quadratures can be applied to the integration of any function f together with the weight function as

$$\int_{\Omega} f(x)\, w(x)\, dx \approx \sum_{g=1}^{N} f(x_g)\, w_g, \tag{1.47}$$

where $\{x_g\}$ are quadrature points and $\{w_g\}$ are the corresponding quadrature weights. The quadrature points are not necessarily equidistant and the weights are not necessarily equally valued, in contrast to other integration schemes such as the trapezoidal rule.

This approximation is exact if the integrand f can be approximated in terms of a linear combination of the polynomials up to an order smaller than $2N-1$. This property is a major benefit of Gaussian quadratures compared to other methods: the exactitude of the approximation is achieved for a fixed number of points, instead of in the limit of infinitely many points. However, this also introduces limitations, as the approximation of the integral may be innacurate if this threshold is not surpassed. For a detailed treatment of Gaussian quadrature theory and construction of quadrature rules, we refer the reader to Davis and Rabinowitz [28].

Using Gaussian quadratures, the integrals of the overlap of functions in the basis set is approximated by

$$\int_{\Omega} \phi_i(x)\phi_j(x)\, dx = \int_{\Omega} p_i(x)p_j(x)w(x)\, dx \approx \sum_{g=1}^{N} p_i(x_g)p_j(x_g)w_g = \delta_{ij}. \tag{1.48}$$

Because $p_i(x)p_j(x)$ generates a polynomial of order $i+j$, the approximation of this integral is exact if $2N > i+j+1$.

In quantum physics, the required integrals usually take the form $\langle\psi_i|O|\psi_j\rangle$, where O is an operator. For function-valued operator, such as the PES $V(x)$, the integral can be computed using Gaussian quadratures as

$$\langle\psi_i|V|\psi_j\rangle = \int_{\Omega} dx\; p_i(x)V(x)p_j(x)w(x) \approx \sum_{g=1}^{N} p_i(x_g)p_j(x_g)V(x_g)w_g. \tag{1.49}$$

The total number of points for finding an exact approximation depends on the degree of the polynomial that is needed to approximate $p_i(x_g)p_j(x_g)V(x_g)$, and, therefore, on the properties of the PES. For PESs with quickly diverging values or lacks of smoothness, a very large number of quadrature points may be necessary.

Gaussian quadrature theory is inheritly a univariate theory. To extend it to d-dimensional spaces, a direct-product quadrature grid is constructed by taking tensor products of univariate quadrature points. The quadrature weight of each multidimensional point is the multiplication of the weights in each dimension. This structure naturally aligns with direct-product basis sets. If n quadrature points are used per dimension, the total number of points is n^d. Recalling that the number of points needed for high accuracy roughly scales with the number of functions in the truncated set, the total number of points scale exponentially with dimensionality— another expression of the curse of dimensionality.

Despite the exponential scaling, in practice, the use of quadratures represents a soft bottleneck of the curse of dimensionality, meaning that it drastically affects the efficiency of the calculation but does not make the computation impossible. Given that the integration result is a sum, it does not require to store the evaluation at each point. Therefore, the calculation can be divided in batches of points that can be processed sequentially, summing their contributions. This allows for memory-efficient integration. As long as the resulting matrix can be stored in memory, the evaluation of the integral remains tractable, though overall computation time and efficiency is impacted.

1.5.2 Monte Carlo integration: importance sampling

In contrast to Gaussian quadratures, the convergence of Monte Carlo integration with respect to the number of points is independent to the dimensionality of the problem. For this reason, Monte Carlo is typically selected as the integrator for high-dimensional systems, i. e., of dimension bigger or equal to 6. In particular, importance sampling Monte Carlo is well-fitted for the required integrals in quantum physics, as $|\psi|^2$ is a probability density function that can be used for sampling.

We focus on the application to the integrals in quantum physics which are broadly given in the form $\langle\psi_i|O|\psi_j\rangle$, where O is an operator. A common choice of importance sampling Monte Carlo that approximates this integral is

$$\langle\psi_i|O|\psi_j\rangle = \int_\Omega d\mathbf{x}\, |\psi_i|^2 \frac{O\psi_j}{\psi_i} = \mathbb{E}_{\mathbf{x}\sim|\psi_i|^2}\left(\frac{O\psi_j}{\psi_i}\right) \approx \frac{1}{N}\sum_{\mathbf{x}_g\sim|\psi_i|^2}^{N} \frac{O\psi_j}{\psi_i}(\mathbf{x}_g), \tag{1.50}$$

where $\mathbb{E}(\cdot)$ is the expected value of the function according to a density sampling and $\mathbf{x}_g \sim |\psi_i|^2$ denotes that the points are drawn from the probability density $\rho_i = |\psi_i|^2$. The step of creating the points according to the probability density ρ_i is called importance sampling. Notably, due to this way of creating points, no point will be sampled in the regions where $\rho_i = 0$, which makes the evaluation of $\frac{O\psi_j}{\psi_i}$ always numerically stable.

Under this scheme, the approximation of the matrix representation of the operator O_{ij} requires I different sets of points, one per density ρ_i. If N points are used in each dimension, a total of $N \cdot I$ points is required. This provides a linear scaling with the number of functions in the truncated set, compared to the exponential dependence in Gaussian quadratures.

The quantity of interest for the algorithm is the integral, which is created as a sum over points. For this reason, the values at individual points are not required to be stored in memory. Dividing the set of points into batches that can be fitted in memory makes Monte Carlo integration a soft bottleneck for the computation of the Hamiltonian matrix.

Alternatively, a Monte Carlo integral using importance sampling can be created without using the quantities in the integrand. By introducing a probability density ρ, the integral $\langle\psi_i|O|\psi_j\rangle$ can be approximated as

$$\langle\psi_i|O|\psi_j\rangle = \int_\Omega d\mathbf{x}\, \rho\, \frac{\psi_i^* O\psi_j}{\rho} = \mathbb{E}_{\mathbf{x}\sim\rho}\left(\frac{\psi_i^* O\psi_j}{\rho}\right) \approx \frac{1}{N}\sum_{\mathbf{x}_g\sim\rho}^{N} \frac{\psi_i^* O\psi_j}{\rho}(\mathbf{x}_g). \tag{1.51}$$

In contrast with the previous Monte Carlo integrator, this approximation can be performed for all the entries of the matrix O_{ij} using a single set of N points, which is now totally independent to the number of functions in the representation.

We provide a brief discussion on how to choose an optimal distribution ρ for sampling. For the integral of the arbitrary function f, the distribution ρ can be used to approximate the integral by

$$I = \int d\mathbf{x}\, f \approx \frac{1}{N} \sum_{x_g \sim \rho}^{N} \frac{f}{\rho}.$$

Assuming the image of f is positive, a possible choice of ρ is $\rho = C\, f$, with $C > 0$ being a normalisation constant. Under this choice, the normalisation requirements of ρ enforces $C = 1/I$. Then, for any number of sampled points N, the approximation of the integral is exact

$$\frac{1}{N} \sum_{\mathbf{x}_g \sim C\, f}^{N} \left(\frac{f}{C\, f} \right) (\mathbf{x}_g) = \frac{1}{C} = I.$$

This constant proportionality is in general not possible, as in principle integrands are not positive everywhere. In addition, the normalization constant C requires the evaluation of the integral I, which is the quanitity that we are trying to approximate.

The difference between the value of the integral I and the approximated value by the Monte Carlo estimator $\hat{I}$ is proportional to the standard deviation of the remaining integrand [29]

$$|I - \hat{I}| \propto \sigma \left(\frac{f}{\rho} \right). \tag{1.52}$$

Using the previous condition $\rho = C\, f$, the obtained variance is zero and the choice of density is optimal. If f is not positive everywhere, a zero variance cannot be obtained, as the division by $\rho = C\, |f|$ will provide both positive and negative numbers. For a family of probability densities ρ_α parametrised by the set α, the optimal choice of parameters α_0 satisfies

$$\sigma \left(\frac{f}{\rho_{\alpha_0}} \right) = \min_{\alpha} \sigma \left(\frac{f}{\rho_\alpha} \right),$$

and therefore an optimization of the probability distribution for sampling can be set up.

In practice, when computing operator matrices O_{ij}, the integrands vary across entries, making the construction of a universally optimal distribution impossible. Instead, the mean of the variance of each element might be chosen as the target of the optimization. As the Hamiltonian matrix is divided in two (or more) operators, a different probability distribution for sampling can be chosen for each of the respective integrals.

Regardless of the approach, the variance of the resulting function cannot be set to zero for all the required integrals in O_{ij}. Then, we study the convergence of the integral with the number of points N. This error of the Monte Carlo integration approximation is governed by the central limit theorem [30], yielding

$$|I - \hat{I}| \sim \frac{1}{\sqrt{N}}. \tag{1.53}$$

This convergence is independent of the dimensionality of the integral, being the main strenght of Monte Carlo methods over Gaussian quadratures. However, it also implies that high-accuracy approximations require $N \to \infty$, and that the rate of improvement is diminished as the number of points increases. This is the downside of Monte Carlo integration schemes with respect to Gaussian quadratures, which provide an approximation that is exact up to numeric precision by surpassing a threshold number of required quadrature points.

2 Enhanced basis sets via normalizing flows[1]

The benefits and difficulties of using a basis set for solving the TISE were layed out in section 1.3.1. The structure of orthonormal basis sets aligns well with the orthonormal structure of the solutions of the TISE. On the other hand, constructing basis sets and calculating the Hamiltonian matrix in the representation of the basis is a computational demanding step. This challenge is aggravated in high-dimensional calculations— this is the curse of dimensionality.

A sensible strategy to improve the methodology to solve the TISE is to build upon spectral methods, conserving the good structural properties while targeting to enhance the convergence of the method to the solution.

In this chapter, we introduce a methodology for the construction of flexible basis sets that can be optimized for the system at hand. These basis sets are constructed using an underlying basis set. Then, a change of coordinates that maps the domain of the basis set to the system domain is applied. The change of coordinates is parametrized by a normalizing flow, i. e., an invertible neural network. If the underlying basis set was orthonormal, the algorithm preserves its orthonormalility while reducing the required number of functions per dimension needed to obtain high-accuracy. In this case, the underlying basis set is built as a direct-product basis set of univariate functions, although this is not a requirement of the methodology.

2.1 Introduction

Accurate calculations of highly excited vibrational states of polyatomic molecules are essential for unravelling increasingly rich experimental spectroscopic information and understanding the dynamics of intermolecular motions. The highly excited molecular vibrations are especially important in fields such as chemical reactivity [31–33] and collisions [34, 35], relaxation processes [36] and stimulated emission [37], as well as spectroscopic probing of high-temperature environments found on exoplanets [38, 39] and in industrial applications [40, 41].

A range of variational and perturbative methods were developed for predicting vibrational spectra of molecules [16, 42–47]. These methods solve the eigenvalue problem for a vibrational Hamiltonian, which is constructed using appropriately chosen vibrational coordinates. The choice of coordinates is a crucial task that directly affects the accuracy of energy calculations. When using a direct product basis of univariate functions, a key challenge is selecting coordinates that provide a large degree of separability of vibrational motions, thereby reducing the computational effort required to solve the vibrational eigenvalue problem [48]. This is particularly important for calculations of delocalized vibrational states of floppy molecules [49, 50], such as van der Waals complexes [51], molecules near dissociation [52], or high-energy excitations in general [53], where couplings between different vibrational modes are prominent.

Rectilinear normal coordinates provide a natural starting point for seeking separability in vibrational problems. However, they become less effective for highly excited states

---

[1] This chapter is based on the publication by Yahya Saleh, Álvaro Fernández Corral, Emil Vogt, Armin Iske, Jochen Küpper and Andrey Yachmenev, ”Computing Excited States of Molecules Using Normalizing Flows”, `https://doi.org/10.1021/acs.jctc.5c00590`, published in the Journal of Chemical Theory and Computation (2025), **21**, 10, 5221–5229. In this work, I contributed to the developement of the theory, production of the code, conduction of calculations, interpretation of results, and writing of the manuscript.

and are generally not suited for floppy molecules, e. g., weakly bound complexes, which naturally sample configurations far from their reference equilibrium geometry. Alternative curvilinear coordinate systems, such as Radau [23], Jacobi [54, 55], valence [56], and polyspherical [57–59] coordinates, were successfully applied in the vibrational calculations of various floppy polyatomic molecules [43, 60]. Choosing the optimal coordinates requires a combination of intuition, consideration of the symmetries of the system, and prior knowledge of the potential energy landscape. This task is particularly challenging for floppy molecules and generally large systems. Several general strategies were recently developed to guide the selection and design of vibrational coordinates, drawing from the available pool of known curvilinear and rectilinear coordinates [61–63].

Due to the diversity of nuclear motions and their dependence on molecular size and bonding topology, no single coordinate system is universally optimal for describing the vibrations of different molecules. One promising approach to improve the effectiveness of a coordinate system involves developing general coordinates, parametrized by variables that can be optimized to minimize vibrational couplings or energy levels. Such general coordinates, expressed as linear combinations of normal coordinates [64, 65], curvilinear coordinates [27, 66–71], or as a quadratic function of normal coordinates [72], were shown to significantly enhance the accuracy of variational calculations. Despite these developments, the broader application of coordinate optimization in variational calculations remains largely unexplored, with previous efforts generally limited to linear parametrizations specific to particular systems.

In this work, we introduce a new general nonlinear parametrization for vibrational coordinates that is based on normalizing flows [73], implemented using a neural network. The parameters of the neural network are optimized using the variational principle. Applied to the calculation of the 100 lowest vibrational states in H_2S, H_2CO, and HCN/HNC molecules, the present approach achieves several orders of magnitude greater accuracy in energy predictions compared with commonly used curvilinear coordinates for the same number of basis functions. The optimized vibrational coordinates effectively capture the underlying physics of the problem, reduce couplings between different vibrational modes, and remain consistent across various levels of basis-set truncation. Building on this property, we propose a cost-efficient approach in which the coordinates are first optimized using a small number of basis functions and then applied to calculations with a larger number of basis functions, keeping the parameters fixed.

2.2 Methods

2.2.1 Enhancing a basis by change of coordinates

We begin by choosing a truncated set of orthonormal basis functions $\{\phi_n(\mathbf{r})\}_{n=0}^{N}$ of L^2 along with an invertible map g_θ parametrized by a set of parameters θ. g_θ maps an initial set of vibrational coordinates $\mathbf{r}$ to a new set of coordinates $\mathbf{q}$ of the same dimension, i. e., $\mathbf{q} = g_\theta(\mathbf{r})$. To improve the approximation properties of the basis functions ϕ_n, we evaluate them in $\mathbf{q}$ to obtain a new set of augmented basis functions $\{\gamma_n(\mathbf{q};\theta)\}_{n=0}^{N}$ defined as

$$\gamma_n(\mathbf{q};\theta) := \phi_n\big(g_\theta(\mathbf{r})\big)\sqrt{|\det \nabla_{\mathbf{r}} g_\theta(\mathbf{r})|}, \tag{2.1}$$

where multiplying by the inverse of the square root of the determinant of the Jacobian ensures that the basis functions remain orthonormal with respect to the L^2-inner product, independent of the values of θ. Inducing augmented basis functions by a nonlinear change of variables is analogous to inducing an augmented probability distribution p from a base distribution p_0 using a change of variables g_θ, commonly referred to as a normalizing flow in the machine learning literature [73, 74]. Therefore, we refer to g_θ as a normalizing flow and to $\mathbf{q}$ as normalizing-flow coordinates.

The map g_θ can, in principle, be any differentiable invertible function. However, to maintain the completeness of the augmented basis set $\{\gamma_n(\mathbf{q};\theta)\}_{n=0}^{\infty}$, the normalizing flow must be bijective [75]. We construct g_θ using an invertible residual neural network (iResNet) [76]. We refer to the supplementary information for a more detailed explanation of the equivalence between optimizing basis sets and vibrational coordinates.

2.3 Architecture

By construction, the iResNet, which is commonly used for image processing, places no restrictions on the image domain. However, in many computational physics and chemistry applications, the domain of internal coordinates is inherently bounded. For example, in vibrational calculations, internal coordinates often represent distances, which are strictly positive, or angles, which are typically periodic or confined to a finite interval such as $[0,\pi]$ or $[0,2\pi]$. Mapping into non-physical or redundant ranges of internal coordinates can lead to inaccurate numerical outcomes and violations of the variational limit. For example, the iResNet output may extend into coordinate regions where the potential is undefined, leading to unreliable results.

To address these issues, we developed an invertible flow that enables control over the output ranges. First, we mapped the ranges of the quadratures of the basis sets in each dimension to $[-1,1]$ using a fixed linear scaling. This way, we allow any domain of the primitive basis set to be handled identically. Next, the iResNet was applied, producing an unbounded output. To map the output back to a finite interval, we applied a wrapper function, L. This wrapper function can be any mapping that transforms an infinite domain into a finite domain. We selected the $\tanh$ function to map values to the interval $[-1,1]$. Finally, we applied a linear scaling $\mathbf{a}\cdot\mathbf{x}+\mathbf{b}$ to adjust the output to the desired target interval. The scaling parameters $\mathbf{a}$ and $\mathbf{b}$ were optimized as part of the workflow. To ensure the output remains within the correct interval, we used a subparametrization of the linear parameters $a_i(\alpha)$, $b_i(\beta)$, where α and β are optimizable variables. We applied different functions for different coordinates, depending on their specific range requirements. Three different subparametrizations were used to accommodate infinite, semi-infinite, and finite intervals. A schematic of this workflow is provided in Fig. 2.1.

2.3.1 Construction of the Hamiltonian

The matrix elements of the vibrational kinetic and potential energy operators in the augmented basis (2.1) can be expressed by introducing a change of variables $\mathbf{q}=g_\theta(\mathbf{r})$ in the integrals. For the potential, this results in the following expression

$$\mathbf{V}_{n'n}=\langle\gamma_{n'}|V|\gamma_n\rangle=\int\phi_{n'}^{*}(\mathbf{q})V(g_\theta^{-1}(\mathbf{q}))\phi_n(\mathbf{q})\mathrm{d}\mathbf{q}. \tag{2.2}$$

This illustrates that the normalizing flow g_θ effectively modifies the coordinates in which the operators are expressed for a fixed set of basis functions $\{\phi_n\}_n$.

The matrix elements of the kinetic energy operator in the augmented basis set are given by

$$
\begin{aligned}
T_{nn'} &= \int \phi_n^*(\mathbf{q})\hat{T}(g_\theta^{-1}(\mathbf{q}))\phi_{n'}(\boldsymbol{q})\mathrm{d}\mathbf{q} \qquad (2.3)\\
&= \frac{\hbar^2}{2}\sum_{kl}\int \left[\left(\frac{1}{2\sqrt{D}}\frac{\partial D}{\partial q_k} + \sqrt{D}\frac{\partial}{\partial q_k}\right)\phi_n^*(\mathbf{q})\right]\\
&\times \sum_{\lambda\mu}\frac{\partial q_k}{\partial r_\lambda}G_{\lambda\mu}(g_\theta^{-1}(\mathbf{q}))\frac{\partial q_l}{\partial r_\mu}\\
&\times \left[\left(\frac{1}{2\sqrt{D}}\frac{\partial D}{\partial q_l} + \sqrt{D}\frac{\partial}{\partial q_l}\right)\phi_{n'}(\mathbf{q})\right]\mathrm{d}\mathbf{q}.
\end{aligned}
$$

$D = \left|1/\det \nabla_{\mathbf{q}} g_\theta^{-1}(\mathbf{q})\right|$, $G_{\lambda\mu}$ is the kinetic-energy matrix, λ, μ indices are used to denote the elements of the coordinate vector $\mathbf{r}$, and k, l indices denote the elements of the coordinate vector $\mathbf{q}$, i. e., $q_k = g_{\theta,k}(\mathbf{r})$ and $r_\lambda = g_{\theta,\lambda}^{-1}(\mathbf{q})$. The differential operators, $\frac{\partial}{\partial q_k}$, only operate inside the square brackets. To obtain this formula, we employed integration by parts, enabling the second-order derivative operator to act symmetrically on both the bra and ket functions as first-order derivatives. The boundary term is omitted as its contribution is zero for most integration domains. Additionally, the kinetic energy operator includes the so-called pseudopotential term, which originates from the transformation from Cartesian to the initial internal coordinates. It is calculated as

$$
U = \frac{\hbar^2}{32}\sum_\lambda\sum_\mu \frac{G_{\lambda\mu}}{\tilde{g}^2}\frac{\partial \tilde{g}}{\partial r_\lambda}\frac{\partial \tilde{g}}{\partial r_\mu} + 4\frac{\partial}{\partial r_\lambda}\left(\frac{G_{\lambda\mu}}{\tilde{g}}\frac{\partial \tilde{g}}{\partial r_\mu}\right), \qquad (2.4)
$$

where $\tilde{g} = \det\left(G^{-1}\right)$. The pseudopotential is a scalar operator and its matrix elements in the transformed basis can be expressed analogously to the potential energy matrix elements in (2.2).

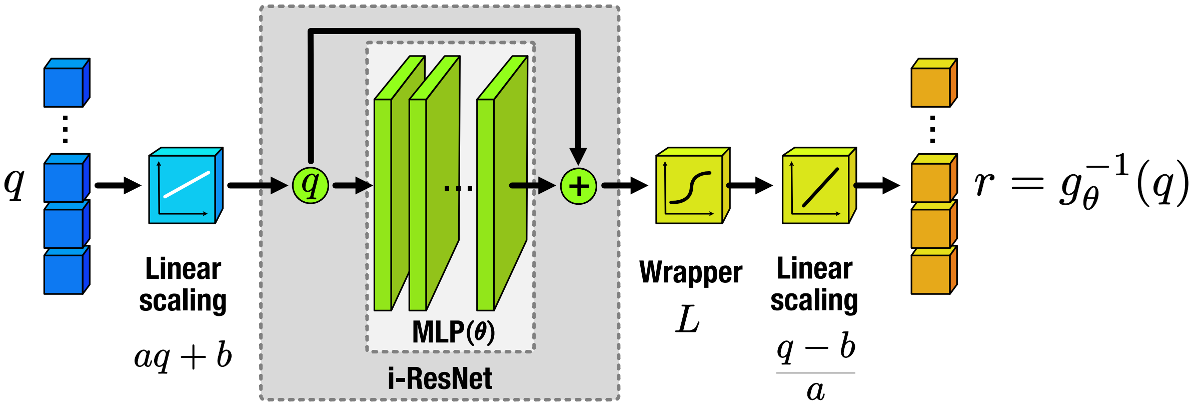

Figure 2.1: **Schematic diagram of the computational workflow of the normalizing-flow function.** As wrapping functions we used $L = \tanh(x)$. A fixed scaling procedure was applied to map the initial coordinate ranges to an identical domain. The linear scaling parameters a and b were optimized together with the MLP θ-parameters.

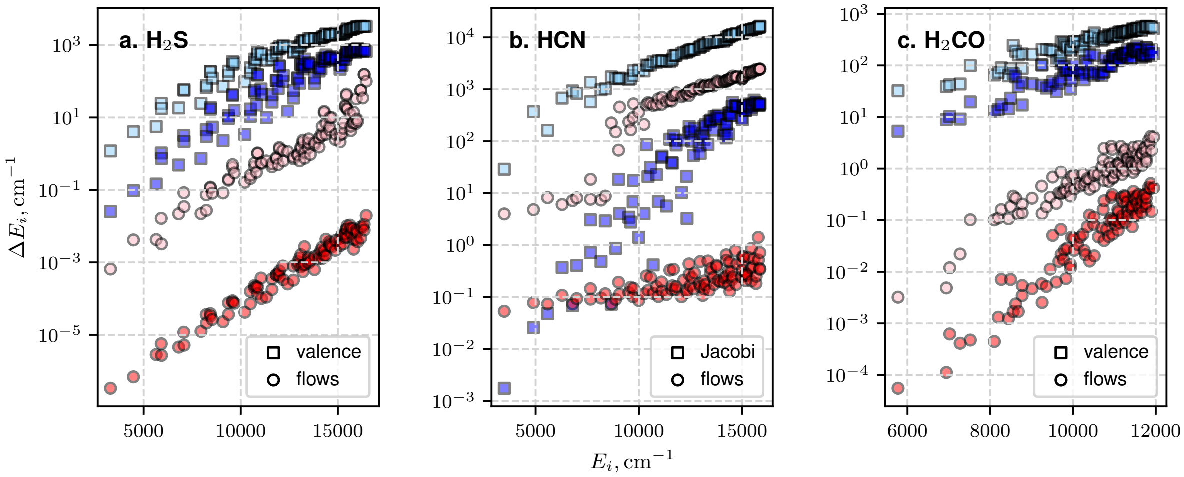

Figure 2.2: **Convergence of H_2S, HCN and H_2CO vibrational energy levels.** Plotted are the 100 lowest energy discrepancies using standard (blue squares) and normalizing-flows (red circles) coordinates. Light and dark colors represent truncations corresponding to a smaller and larger number of basis functions, respectively. **a** Energy discrepancies for H_2S at $P_{\max} = 12$ (140 basis functions) and 20 (506). **b** Energy discrepancies for HCN at $P_{\max} = 12$ (140) and 32 (1785). **c** Energy discrepancies for H_2CO at $P_{\max} = 9$ (1176) and 12 (3906).

2.3.2 Optimization

We approximate the vibrational wavefunctions Ψ_m ($m = 1..M$) as linear combination of augmented basis functions (2.1), i. e.,

$$\Psi_m(\mathbf{q}) \approx \sum_{n \leq N} c_{nm} \gamma_n(\mathbf{q}; \theta). \tag{2.5}$$

The linear-expansion coefficients c_{nm} and the normalizing-flow parameters θ are determined using the variational principle, by minimizing the energies of the ground and excited vibrational states. For the coefficients, this is equivalent to solving the eigenvalue problem $\mathbf{E} = \mathbf{C}^{-1}\mathbf{H}\mathbf{C}$, where $\mathbf{C} = \{c_{nm}\}_{n,m}^{N,M}$, $\mathbf{E} = \{E_m\}_m^M$ are the vibrational energies, and $\mathbf{H} = \mathbf{T} + \mathbf{V}$ is the sum of matrix representations of the kinetic and potential energy operators, given by (2.2) and (2.3).

Because vibrational energies are nonlinear functions of the parameters θ, these parameters are optimized using gradient descent methods. The optimization is guided by a loss function derived from the variational principle and may involve minimizing quantities such as the sum of vibrational energies, the trace of the Hamiltonian matrix, or the matrix exponential. A loss function expressed as the sum of all energies spanned by the chosen basis set is equivalent to the trace of the Hamiltonian matrix,

$$\mathcal{L}_\theta = \sum_{n \leq N} E_n = \mathrm{Tr}(\mathbf{H}) \rightarrow \min_\theta . \tag{2.6}$$

This loss function has a relatively low computational cost, as it decouples the nonlinear parameters θ from the eigenvector coefficients c_{nm}, requiring only the evaluation of diagonal elements of the Hamiltonian matrix when the initial basis is orthonormal. In contrast, when the loss function is based on the sum of a subset of the lowest energies, the parameters depend on the eigenvector coefficients c_{nm}, and a repeated solution of the eigenvalue problem during optimization is required. Despite the added complexity, the high accuracy achieved even with small number of basis functions can potentially outweigh the computational costs of the repeated matrix diagonalization. In our calculations we used the sum of a subset of all vibrational energies as the loss function, *vide infra*. A cost-efficient optimization and application strategy that mitigates the cost of repeated matrix diagonalization is discussed in subsection 2.5.4.

The evaluation of the matrix elements in (2.2) and (2.3) is one of the most computationally demanding parts. In this work, we employed Gaussian quadratures to compute the necessary integrals, altering the quadrature degree at different optimization steps to prevent overfitting. We found that alternating between smaller quadratures during optimization was computationally more efficient while still converging to the same values of the parameters θ as those obtained using a larger quadrature. After convergence, the final energies and wavefunctions were computed by solving the eigenvalue problem with a large quadrature for accurate integral evaluations. For higher-dimensional systems, more efficient techniques such as sparse-grid methods [77] or collocation [78] can be used. Alternatively, Monte-Carlo methods [79, 80] may be employed when high accuracy is not required.

2.4 Computational details

The accuracy and performance of our approach were validated in calculations of vibrational states for hydrogen sulfide H_2S, formaldehyde H_2CO, and hydrogen cyanide/hydrogen isocyanide HCN/HNC isomers. For H_2S and H_2CO, we used valence coordinates as the reference coordinates and employed a direct product of Hermite functions as the basis set. For H_2S, the direct product basis was constructed by considering only combinations of one-dimensional vibrational quantum numbers (n_1, n_2, n_3) that satisfy the polyad condition $2n_1 + 2n_2 + n_3 \leq P_{\text{max}}$, where n_1, n_2, and n_3 correspond to the vibrational quanta for r_{SH_1}, r_{SH_2}, and $\alpha_{\angle\text{H}_1\text{SH}_2}$ valence coordinates, respectively. For H_2CO, we applied the basis truncation condition $2n_1 + 2n_2 + 2n_3 + n_4 + n_5 + n_6 \leq P_{\text{max}}$, where $n_1, \ldots, n_6$ correspond to the vibrational quanta for valence coordinates r_{CO}, r_{CH_1}, r_{CH_2}, $\alpha_{\angle\text{OCH}_1}$, $\alpha_{\angle\text{OCH}_2}$, and τ, the dihedral angle between the OCH_1 and OCH_2 planes. For HCN/HNC, we used the basis truncation condition $2n_1 + 2n_2 + n_3 \leq P_{\text{max}}$ and Jacobi reference coordinates $r_{\text{CN}}, R, \alpha_{\angle\text{R-CN}}$, where R is the distance between the hydrogen atom and the center of mass of the C–N bond, and $\alpha_{\angle\text{R-CN}}$ is the angle between these coordinate vectors. Hermite functions were used for the two radial coordinates and Legendre functions for the angular coordinate. The Legendre functions were multiplied by $\sin^{1/2}(\alpha_{\angle\text{R-CN}})$ to ensure the correct behaviour of the wavefunction at the linear geometry of the molecule, where the Hamiltonian becomes singular [81, 82]. For all molecules, we employed spectroscopically refined potential energy surfaces (PES) [83–85] and numerically constructed exact kinetic energy operator using the method described in [43, 53].

To model the normalizing flow g_θ, we used an iResNet consisting of 10 blocks. Each block was represented by a dense neural network comprising two hidden layers with unit sizes $[8, 8]$ and an output layer of n units, where n corresponds to the number of coordinates. A more detailed description of the normalizing-flow architecture is available in the supplementary information. The normalizing-flow parameters were optimized variationally by minimizing the sum of the 100 or 200 lowest vibrational energies. Generally, 1000 iterations were enough to achieve good convergence. Benchmark energies were computed with basis sets truncated at $P_{\text{max}} = 60$ (optimized for $P_{\text{max}} =$

12) for H_2S, $P_{\max} = 16$ (optimized for $P_{\max} = 9$) for H_2CO, and $P_{\max} = 44$ for HCN/HNC.

2.5 Results

2.5.1 Computed Vibrational Energies

On average, over the 100 lowest energies, the calculations converged with an accuracy of 0.04 cm^{-1} for H_2S, 0.53 cm^{-1} for H_2CO, and 0.03 cm^{-1} for HCN/HNC compared to the reference values reported in the literature [83–85]. We thus considered our results to be converged and used them as benchmark data throughout the rest of the manuscript. A table summarizing the deviations of vibrational energies from the reference values is provided in the supplementary information.

The absolute error for the 100 lowest vibrational states of H_2S, H_2CO, and HCN/HNC, as a function of the basis-set truncation parameter $P_{\max}$, is shown in Fig. 2.2. For each molecule, the results of two variational calculations are presented, one using reference valence or Jacobi coordinates and another using the optimized normalizing-flow coordinates. With the same number of basis functions, coordinate optimization resulted in up to five orders of magnitude improvement in the accuracy of vibrational energy calculations compared to using the standard reference coordinates. Extrapolating to a larger number of basis functions, we estimate that matching the same accuracy using the reference coordinates would require approximately an order of magnitude increase in the number of basis functions.

We note that the convergence of results can also be improved by increasing the complexity of the normalizing-flow function. A detailed analysis of this effect, along with an investigation of how varying the number of target states affect the normalizing-flow coordinates, is provided in the supplementary information.

Direct product basis sets can be improved using basis-set contraction, which involves partitioning the total Hamiltonian into subsystems, solving reduced-dimensional variational problems for each, and then using these solutions for the full-dimensional problem [22, 86]. For molecules like H_2S and H_2CO, basis set contraction works well due to near-separability of valence coordinates in the PES. However, this approach becomes more challenging for floppy molecules like HCN/HNC, where two of the Jacobi coordinates are strongly coupled.

This challenge is illustrated in Fig. 2.3 for the first 200 vibrational energies of HCN/HNC. The figure presents results of basis set contraction using Jacobi coordinates alongside those obtained with a direct product basis set of primitive functions, i. e., Hermite and Legendre, using optimized normalizing-flow coordinates. The contracted basis was constructed by partitioning the Hamiltonian into one-dimensional subsystems for each of the Jacobi coordinates, with the HCN isomer equilibrium geometry as the reference configuration. As shown in Fig. 2.3, the contracted basis significantly improves convergence compared to the product basis set results in Fig. 2.2 b, but not for all vibrational states. High accuracy is only achieved for states localized around the HCN minimum of the PES, while delocalized states and states localized around the HNC minimum show little improvement. In contrast, the optimized normalizing-flow coordinates provide a balanced description of all localized and delocalized states, with a much smaller spread in errors across different states.

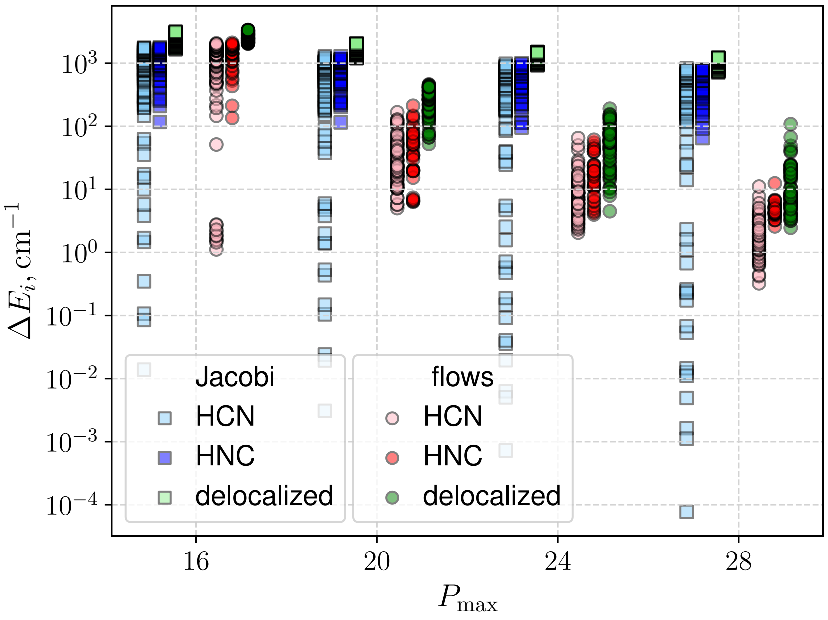

Figure 2.3: **Convergence of HCN/HNC vibrational energy levels.** Shown are the lowest 200 energies for using Jacobi coordinates (squares) and normalizing-flow coordinates (circles). The energy discrepancies (ΔE_i) relative to our converged benchmark reference are shown for several basis sets, truncated at $P_{\max} = 12$ (140 basis functions), 16 (285), 20 (506), 24 (811), and 28 (1200). Vibrational states assigned to the HCN isomer, the HNC isomer, and states with an energy above the isomerization barrier (delocalized) are differentiated by color. All states are slightly offset along the $P_{\max}$ axis for visual clarity.

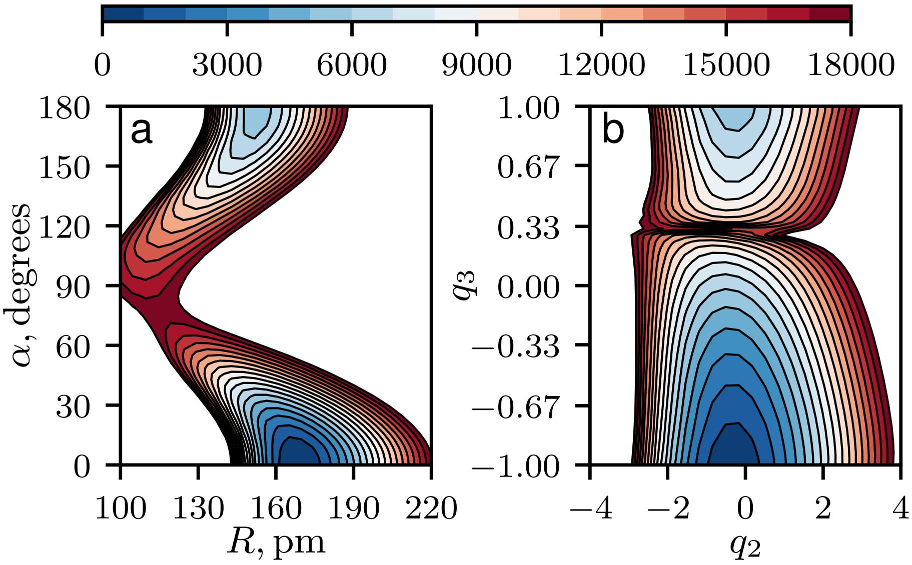

Figure 2.4: **Two-dimensional cuts of the HCN/HNC potential energy surface. a** Cut along the Jacobi coordinates R and $\alpha_{\angle\text{R-CN}}$. **b** Cut along the optimized normalizing-flow coordinates. The optimization was performed for a basis set truncated at $P_{\max} = 16$ (285 basis functions) with the loss function defined as the sum of the 100 lowest vibrational state energies. An effective decoupling of the surface in the normalizing-flows coordinates is patent.

2.5.2 Intepretability

To gain insight into the interpretation of the optimized normalizing-flow coordinates, we plotted in Fig. 2.4 the two-dimensional cut of the PES of HCN/HNC along the strongly coupled Jacobi coordinates, R and $\alpha_{\angle\text{R-CN}}$, i. e., $V(r_0, R, \alpha_{\angle\text{R-CN}})$, alongside the corresponding cut in the optimized normalizing-flow coordinates, $V(g_\theta^{-1}(q_1, q_2, q_3))$. In these plots, r_{CN} coordinate is fixed at its equilibrium value r_0 and $q_1 = 0$. The potential is clearly highly anisotropic when expressed in Jacobi coordinates (panel a), which leads to a strong coupling between the two vibrations. In contrast, when expressed in the optimized coordinates (panel b), the HCN↔HNC minimum energy isomerization pathway is practically a straight line along the coordinate q_3 at $q_2 \approx 0$. This reduction in anisotropy explains why coordinate optimization improves the convergence in the product basis. The optimization achieves an effective coordinate decoupling of the PES, which allows for a better approximation of the eigenfunctions of the Hamiltonian by the chosen direct product basis. In addition, it is evident from the spacing between the contour lines along the flows coordinate q_2, that the potential becomes more harmonic in this dimension in comparison to R in the Jacobi coordinates. The same behavior was observed when comparing q_1 and r_{CN}. This is expected, as Hermite functions - the solutions of the quantum harmonic oscillator - were used as the basis for stretching coordinates.

2.5.3 Assignment of approximate quantum numbers

Assigning approximate quantum numbers to computed eigenstates connects numerical results to their spectroscopic interpretation. Typically, as the complexity of the method for solving the Schrödinger equation increases, so does the difficulty of assignment. In many cases, less accurate but more interpretable effective models are more practical than highly accurate methods, as they facilitate approximate quantum number assignment and enhance the interpretability of experimental spectra.

The enhanced coordinate decoupling in the HCN/HNC Hamiltonian suggests that assigning approximate quantum numbers to computed eigenstates is more straightforward in normalizing-flow coordinates compared to reference Jacobi coordinates. In Table 1, we compare the accuracy of projection-based assignment of approximate quantum numbers for the first 100 eigenstates of HCN/HNC. Projections were performed onto one-dimensional eigenfunctions (a contracted basis) expressed in either Jacobi or normalizing-flow coordinates.

Measure	Median	Mean	Min	N_{assign}
Jacobi (HCN)	0.27	0.35	0.05	24
Jacobi (HNC)	0.17	0.25	0.04	15
Jacobi (HCN/HNC)	0.39	0.46	0.12	39
Flows (HCN)	0.80	0.79	0.40	97

Table 1: **Projection-based assignment metrics for vibrational states of HCN/HNC**. The tabulated values are the median, mean and minimum values of the largest norm-square projection coefficients obtained by projecting the vibrational wavefunctions for 100 states onto products of one-dimensional eigenfunctions. In the last column, N_{assign}, the number of vibrational states (out of 100) that could be unambiguously assigned a unique set of approximate quantum numbers are shown.

In the limit of convergence of both the three- and one-dimensional eigenfunctions, the projection-based assignment depends only on the choice of coordinates and not on the basis used to compute the eigenfunctions. A unique assignment is ensured when the norm-square of the largest absolute projection coefficient is larger than 0.5. The normalizing-flow coordinates were optimized for 100 eigenstates with $P_{\max} = 28$.

The contracted basis was constructed by partitioning the Hamiltonian into one-dimensional subsystems corresponding to each coordinate. For Jacobi coordinates, either the HCN or HNC equilibrium geometry was used as reference configurations. In contrast, for normalizing-flow coordinates, only the HCN equilibrium geometry was needed due to reduced vibrational coupling. The assignment of approximate quantum numbers is significantly more accurate using the normalizing-flow coordinates with 97 uniquely assigned states compared to only 39 using Jacobi coordinates. This highlights the benefits of optimized coordinate transformations for more reliable spectroscopic interpretations of the results.

2.5.4 Transferability

We found that the converged iResNet parameters remained nearly identical when optimized using different basis-set truncations, suggesting that there exist unique optimal vibrational coordinates for a given type of basis. Leveraging this finding, we developed a cost-efficient approach where the flow coordinates are first optimized using a small number of basis functions and then applied with fixed parameters in calculations with a larger number of basis functions. The size and computational cost of the quantities that depend on the normalizing flow, such as $\frac{\partial q_\alpha}{\partial r_l}$, $\frac{\partial D}{\partial r_l}$, $\frac{1}{D}$, etc., are independent of the number of employed basis functions. This means that calculations using fixed (pre-trained) normalizing-flow parameters scale with the truncation parameter, $P_{\max}$ in the same way as those using a regular linear mapping. The transferability property significantly reduces the computational costs and makes it feasible to apply our approach to high-dimensional systems, while maintaining accuracy comparable to full optimization.

In Fig. 2.5, we show the convergence for the 100 first energy levels of H_2S with respect to the basis-set truncation parameter, $P_{\max}$. The results are obtained using valence and optimized normalizing-flow coordinates. Two types of normalizing-flow coordinates are compared: those optimized for each specific $P_{\max}$ (Opt. flows) and those optimized for selected values of $P_{\max}$ (12, 16, 20) and subsequently transferred to calculations with larger $P_{\max}$. The metric for the convergence is the error of the individual energy levels ($E_i - E_i^{(\mathrm{Ref})}$), where $E_i^{(\mathrm{Ref})}$ represent the benchmark energies detailed in the supplementary information. The results clearly demonstrate that energy calculations using transferred coordinates yield greater accuracy than those using valence coordinates. Moreover, their performance is on par with the more computationally intensive Opt. flows coordinates. The findings also indicate that transferring from a larger $P_{\max}$ can enhance the accuracy of highly excited states.

The convergence of the approximation using the Hermite basis with respect to the number of basis functions, N, is algebraic [2]. Specifically, the error satisfies

$$\left\|\Psi_m - \hat{\Psi}_m\right\| < AN^{-k},$$

where$\|\cdot\|$ denotes the L^2-norm, Ψ_m is the exact wavefunction, and $\hat{\Psi}_m$ is its approximation. The constant A depends on the relationship between the target wavefunction and the

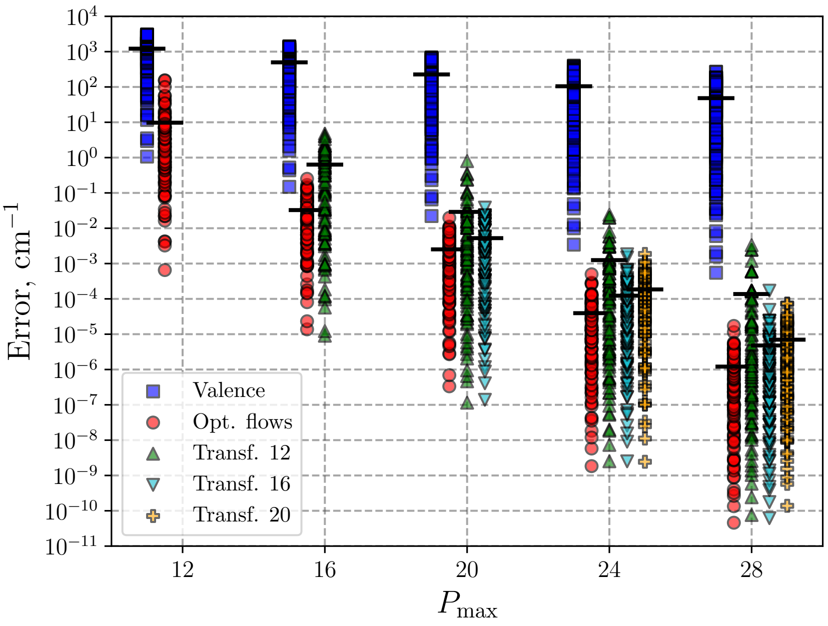

Figure 2.5: **Convergence of the first 100 vibrational energies of H_2S.** Convergence of the first 100 energy levels of H_2S using valence (blue squares), and optimized normalizing-flows (red circles) coordinates, as a function of P_{max}. Results obtained with normalizing-flow coordinates optimized for $P_{\mathrm{max}} = 12$ (green up-triangles), $P_{\mathrm{max}} = 16$ (cyan down-triangles), and $P_{\mathrm{max}} = 20$ (orange plus signs) and applied to calculations with larger P_{max} are also shown. Thick horizontal lines indicate the average energy-level error for each P_{max}, which is minimized through training. Data points are slightly offset along the P_{max} axis for improved visual clarity.

operators associated with the Hermite basis, and k denotes the rate of convergence. The convergence for the Hermite basis defined in normalizing-flow coordinates is also algebraic [87], with different constants A and k for each map. Therefore, it is reasonable to assume that the loss function defined in (3.13) converges algebraically, i. e.,

$$\mathcal{L}(N) = \frac{\mathcal{L}_{\theta}^{100}(N) - \mathcal{L}_{\mathrm{Ref}}^{100}}{100} \sim A\, N^{-k},$$

where θ are the optimized parameters for the chosen N. To quantify the improved convergence rate observed for the normalizing-flow coordinates (see Fig. 2.5), we fitted $\log(\mathcal{L})$ with a linear expression in N, i. e., $\log(\mathcal{L}) = -k\log(N) + \log(A)$. The regression parameters derived from this fit are shown in Table 2. The convergence rate of the two normalizing-flow coordinates is significantly higher than that of the valence coordinates. Remarkably, the convergence rate of the transferred normalizing-flow coordinates reaches 75% of the convergence rate of the flow coordinates optimized at each truncation level. The constant A is also decreased by the use of nonlinear coordinates, which means that the accuracy is improved for any fixed truncation.

The results for $P_{\mathrm{max}} = 12$ for H_2CO in Fig. 2.2 c were calculated with normalizing-flow coordinate optimized for $P_{\mathrm{max}} = 9$. Additional results provided in the supplementary information further demonstrate the utility of the transferability property. In future work, we will elaborate on the transferability property across basis-set truncations and investigate the extension of the principle of transferability of normalizing-flow coordinates to different isotopologues and to molecular systems sharing similar structural motifs. This could potentially contribute to our understanding of intrinsic vibrational coordinates.

Coordinate	$k \times 10^3$	$\log(A)$
Valence	0.61±0.08	11.7±0.3
Opt. flows	2.90±0.47	6.43±2
Transf. 12	2.17±0.30	7.11±1

Table 2: Convergence parameters for different coordinates.

2.6 Conclusion

In summary, we introduced a general nonlinear parametrization for vibrational coordinates of molecules using normalizing flows. By optimizing the flow parameters through the variational principle, we significantly accelerated basis-set convergence, leading to more accurate vibrational energies. The improvement is especially pronounced for highly excited and delocalized vibrational states. The learned coordinates enhanced the separability of the Hamiltonian, which we leveraged to improve assignment of approximate quantum numbers by projection onto direct products of one-dimensional eigenfunctions. The enhanced separability also potentially allows for a more intuitive interpretation of the key motifs in strongly-coupled vibrational dynamics. The transferability of the optimized coordinates across different truncation levels provides a computationally efficient protocol for larger molecular system calculations. As other variational approaches, our method suffers from an exponential growth in the size of the product basis as the number of coordinates increases. This challenge has been effectively addressed in the literature using prescreening techniques that selectively retain only the most relevant basis-product configurations for the states of interest [88, 89]. It should be possible to combine the present normalizing-flow approach with state-specific eigenvalue solvers, where the basis-product configurations are tailored to specific vibrational states. One promising method specifically designed for vibrational solutions is the iterative residuum-based RACE algorithm [90]. After the release of the first arXiv version of this work [91], another group integrated the concept of normalizing flows for basis-set augmentation with Monte-Carlo methods and successfully applied it to high-dimensional systems [92].

We also explored the applicability of the normalizing-flow method for excited electronic states, testing it on single-electron systems such as the hydrogen atom, hydrogen molecular ion, and carbon atom in the single-active electron approximation. Results presented in the supplementary information show a significant improvement in basis-set convergence. This suggests promising potential of normalizing-flow method for electronic structure problems, especially since neural-network-based methods for excited state computations remain challenging [80, 93].

2.7 Supplementary information

Learning an optimized vibrational coordinate system $\mathbf{r} \to g_\theta(\mathbf{r}) = \mathbf{q}$ and learning a new basis set $\{\phi_n(\mathbf{r})\}_{n=0}^{\infty} \to \{\phi_n(\mathbf{q})\sqrt{|1/\det \nabla_{\mathbf{q}} g_\theta^{-1}(\mathbf{q})|}\}_{n=0}^{\infty}$ are two faces of the same coin, as seen, e. g., from performing a change of variable in the matrix representation of the potential

$$\mathbf{V}_{n'n} = \langle \gamma_{n'} | V | \gamma_n \rangle = \int \phi_{n'}^*(\mathbf{q}) V(g_\theta^{-1}(\mathbf{q})) \phi_n(\mathbf{q}) \mathrm{d}\mathbf{q}.$$

In the following, we give a detailed description of the latter perspective, highlighting theoretical foundations and linking to the concept of normalizing flows in machine learning. We also provide detailed information on the numerical simulations conducted in this study, along with additional results that can faciliate the application of this method in other domains.

2.7.1 Mathematical foundations

Composing a basis set $\{\phi_n(\mathbf{r})\}_{n=0}^{\infty}$ with an invertible mapping g yields a sequence of functions $\{\phi_n(\mathbf{q})\sqrt{|1/\det \nabla_{\mathbf{q}} g^{-1}(\mathbf{q})|}\}_{n=0}^{\infty}$. Certain choices of the mapping g can produce sequences of functions with improved approximation properties [94]. A common example is the linear map $g(x) = a\ x + b$, where a and b are chosen to align the potential of the problem with the potential that generates the basis set around the equilibrium geometry of the molecule. There are also examples of nonlinear fixed mappings, such as trigonometric or exponential functions.

However, the use of adaptive nonlinear mappings introduces additional complexity. To ensure convergence of the method, the basis set must retain its completeness after compositing with the mapping g. The concept of composing basis sets with an invertible neural network was first introduced in Cranmer et al. [95]. However, the approach was only applied to simple models, and the rationale for employing invertible neural networks was not fully justified or rigorously explored.

The rationale for using invertible neural networks from an approximation-theory perspective was first analyzed in Saleh et al. [96] and later rigorously established in Saleh and Iske [75], where sufficient and necessary conditions for completeness under perturbation by composition operators were characterized.

Normalizing flows are foundational models in the field of generative machine learning [73], rooted in the probability integral transform initially introduced in Fisher [97]. This theorem states that any continuous probability density p can be converted by a change of variables to a uniform distribution on the interval $[0, 1]$. By imposing a few reasonable conditions to p, e. g., finiteness, it can be shown that this change of variables is bijective. As a result, any pair of probability densities p and q that satisfy these conditions can be related through a change of variables, as both can be bijectively mapped into the uniform distribution.

To illustrate this concept, consider the following example, which is commonly used for its simplicity. Let p be an unknown bimodal probability distribution and consider approximating it by the Gaussian distribution p_0. The Gaussian distribution is a simple distribution characterized by only two parameters, the mean and the variance. However, this approximation of p *via* p_0 is inherently flawed since no combination of the mean and the variance can make the Gaussian p_0 exhibit a bimodal behavior. However, as discussed earlier, we can introduce an invertible change of variables $q = g(x)$ which

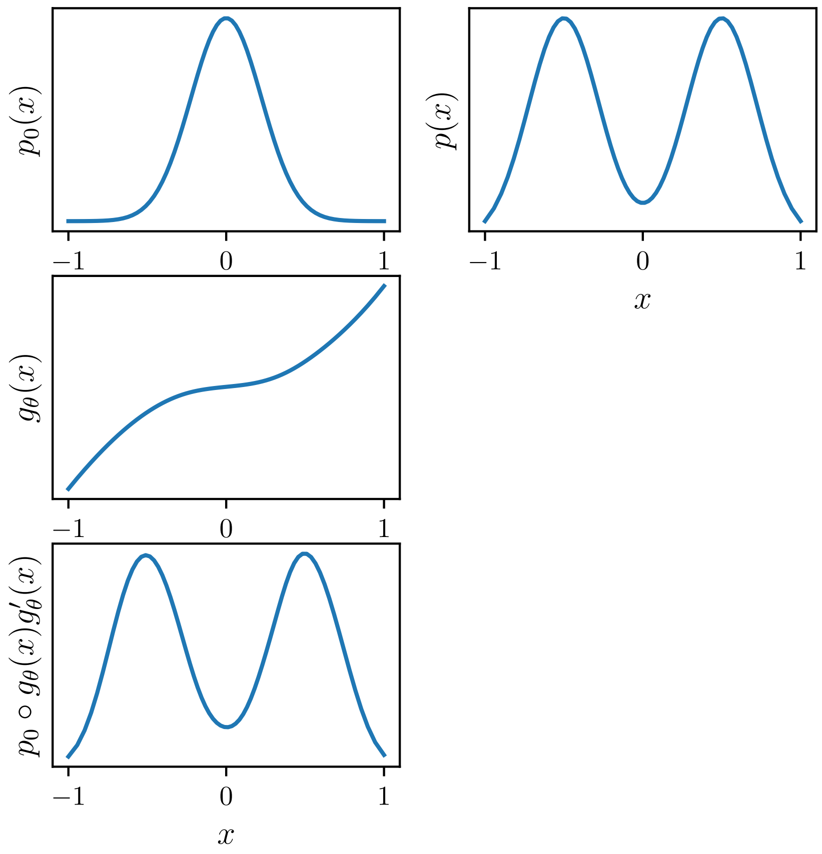

Figure 2.6: An illustration of normalizing flows augmenting a Gaussian distribution p_0. Approximating a bimodal distribution p with a single Gaussian is not possible. However, composing p_0 with the function g_θ produces a good approximation.

allows us to transform the Gaussian distribution p_0 into a more complex distribution p^A that can approximate p. Specifically, we define the new distribution p^A as

$$p^A(x) = p_0(g(x))|\det \nabla_x g(x)|.$$

Multiplying by $|\det \nabla_x g(x)|$, the determinant of the Jacobian of the change of variables, ensures that the distribution remains normalized. Given that both the Gaussian and bimodal distributions are well-behaved, there exists a suitable function g that can bridge this gap. In this case, we parametrize g as an invertible function g_θ, thus augmenting the expressivity of the simple Gaussian model. The results can be found in Fig. 2.6.

If the expressivity of a single Gaussian distribution can be enhanced, why not extend this approach to a family of functions, such as a basis set? This concept forms a core idea of our manuscript.

2.7.2 Computational details

A summary of the atomic masses used in this study are provided in Table 3.

Table 3: The atomic masses utilized in calculations of the vibrational energies for the H_2S, H_2CO and HCN/HNC molecules.

Atom	Mass (u)
Hydrogen	1.00782505
Sulfur	31.97207070
Carbon	12.0
Oxygen	15.99491463
Nitrogen	14.003074004251

Architecture of the invertible residual network

In this work, we model the normalizing flow g_θ as an invertible residual neural network (iResNet) [76]. An iResNet is constructed by concatenating blocks of the form

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{f}_k(\mathbf{x}_k; \theta),$$

where $\mathbf{f}_k$ is the k-th block, $\mathbf{x}_k$ is the input to this block, and $\mathbf{x}_{k+1}$ is the output. We denote the initial set of coordinates as $\mathbf{x}_0$, and for an iResNet with K blocks, we denote its output as $\mathbf{q} = \mathbf{x}_{K-1}$. Each residual block is a standard feed-forward neural network composed of weights, biases and nonlinear activations. Each such block is guaranteed to be invertible if $\mathbf{f}_k$ is Lipschitz, with Lipschitz constant < 1. The inverse of each block can be computed by a fixed-point iteration method.

To guarantee that the blocks $\mathbf{f}_k$ for $k = 0, \ldots, K-1$ satisfy the Lipschitz condition, we used the LipSwish activation function

$$\sigma(x) := \left(\frac{1}{1.1}\right) \frac{1}{1 + \exp(-x)},$$

which has a Lipschitz constant ≈ 1. A block $\mathbf{f}_k$ containing such activation functions is thus guaranteed to be Lipschitz if each of the linear passes has a Lipschitz constant < 1, i. e., the weight matrices W are normalized to have a spectral norm < 1. This is achieved by setting

$$W = \begin{cases} W & \text{if } \|W\|_2 < c \\ c \cdot \frac{W}{\|W\|_2} & \text{if } \|W\|_2 \geq c, \end{cases}$$

where $0 < c < 1$ is a hyperparameter and $\|W\|_2$ is the spectral norm of the matrix W. For a block $\mathbf{f}_k$ containing n hidden layers, it can be shown that

$$\mathrm{Lip}(\mathbf{f}_k) \leq c^{n+1}.$$

A special attention must be given to the architecture of each block. While a block with a Lipschitz constant close to 1 allows for a greater flexibility, a higher Lipschitz constant can significantly hinder the convergence of the fixed-point iteration method used for computing its inverse. To balance these effects, we used blocks with 2 hidden layers for all of our calculations and set $c = 0.9$, producing blocks with Lipschitz constants ~ 0.7.

Comparison against reference data

Benchmark results for vibrational energy calculations were generated using the wavefunction ansatz described in Eq. 3 of the manuscript. The calculations employed a basis polyad truncation $P_{\max} = 60$ for H_2S, $P_{\max} = 16$ for H_2CO, and $P_{\max} = 44$ and for HCN/HCN. We compared our results with those of Azzam et al. [83] for H_2S, Al-Refaie et al. [84] for H_2CO, and Van Mourik et al. [85] for HCN/HNC.

In Table 4, we present the calculated energies E_i, the reference values E_i^{ref} and the deviations $\Delta E_i = E_i - E_i^{\mathrm{ref}}$ across the $i = 1..100$ lowest energy levels for the three molecular systems. The results demonstrate almost exact agreement with the reference data, with most of the states converging to even lower limits than the reference data.

Table 4: Comparison of vibrational energies (in cm^{-1}) of H_2S, HCN/HNC, and H_2CO molecules, calculated using the normalizing-flow approach E_i, with reference energies E_i^{ref} [83–85], across the 100 lowest vibrational states. A positive deviation ΔE_i indicates a better convergence is achieved by our prediction.

i	**H_2S**			**HCN**			**H_2CO**		
	E_i	E_i^{ref}	ΔE_i	E_i	E_i^{ref}	ΔE_i	E_i	E_i^{ref}	ΔE_i
1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2	1182.57	1182.58	-0.01	1414.92	1414.92	-0.00	1167.29	1167.34	-0.05
3	2353.91	2353.96	-0.06	2100.58	2100.58	0.00	1249.06	1249.07	-0.01
4	2614.39	2614.41	-0.01	2801.46	2801.46	-0.00	1500.10	1500.12	-0.03
5	2628.46	2628.45	0.01	3307.74	3307.75	-0.01	1746.03	1746.05	-0.02
6	3513.70	3513.79	-0.09	3510.99	3510.99	0.00	2327.34	2327.50	-0.15
7	3779.19	3779.17	0.02	4176.24	4176.24	0.00	2422.56	2422.63	-0.07
8	3789.27	3789.27	0.00	4181.46	4181.45	0.01	2494.25	2494.32	-0.07
9	4661.61	4661.67	-0.07	4686.28	4686.29	-0.01	2666.81	2667.04	-0.22
10	4932.69	4932.70	-0.01	4891.76	4891.76	0.00	2718.97	2719.08	-0.12
11	4939.13	4939.10	0.03	5185.57	5185.64	-0.07	2782.37	2782.41	-0.04
12	5145.03	5144.99	0.05	5394.43	5394.43	0.00	2843.32	2843.34	-0.01
13	5147.17	5147.22	-0.05	5537.76	5537.76	-0.00	2905.75	2905.86	-0.11
14	5243.16	5243.10	0.06	5586.50	5586.50	0.00	2998.91	2999.01	-0.10
15	5797.21	5797.23	-0.03	6033.72	6033.72	-0.00	2999.96	3000.00	-0.04
16	6074.57	6074.58	-0.02	6127.51	6127.56	-0.05	3238.84	3238.94	-0.10
17	6077.63	6077.59	0.03	6242.43	6242.42	0.00	3471.66	3471.72	-0.06
18	6288.13	6288.15	-0.01	6260.59	6260.59	-0.00	3480.98	3481.30	-0.31
19	6289.13	6289.17	-0.04	6513.48	6513.50	-0.02	3585.63	3585.89	-0.26
20	6385.32	6385.32	-0.00	6768.51	6768.51	-0.00	3675.03	3675.21	-0.18
21	6920.08	6920.08	-0.00	6879.60	6879.60	0.00	3736.63	3737.05	-0.42
22	7204.31	7204.31	-0.00	6960.99	6960.99	-0.00	3825.37	3825.97	-0.60
23	7204.43	7204.44	-0.00	7088.71	7088.74	-0.03	3887.05	3887.36	-0.31
24	7419.85	7419.85	-0.00	7210.51	7210.59	-0.08	3935.85	3936.44	-0.59
25	7420.08	7420.09	-0.02	7369.18	7369.18	-0.00	3941.30	3941.53	-0.23
26	7516.83	7516.83	-0.00	7461.59	7461.60	-0.01	3996.26	3996.48	-0.23
27	7576.41	7576.38	0.03	7617.24	7617.24	0.00	4022.46	4022.56	-0.09
28	7576.60	7576.54	0.05	7641.28	7641.28	0.00	4057.79	4058.10	-0.31
29	7752.34	7752.26	0.08	7855.83	7855.84	-0.01	4083.44	4083.49	-0.05
30	7779.35	7779.32	0.03	8020.41	8020.45	-0.04	4163.89	4164.13	-0.24

31	8029.81	8029.81	-0.00	8110.25	8110.25	-0.00	4164.61	4165.27	-0.66
32	8318.68	8318.68	-0.00	8140.64	8140.69	-0.05	4192.77	4193.19	-0.41
33	8321.86	8321.86	-0.00	8197.55	8197.55	-0.00	4247.40	4247.61	-0.21
34	8539.57	8539.56	0.00	8283.37	8283.37	0.00	4256.12	4256.31	-0.19
35	8539.82	8539.82	-0.00	8323.52	8323.52	-0.00	4335.01	4335.11	-0.10
36	8637.16	8637.16	-0.00	8584.71	8584.74	-0.03	4397.64	4398.13	-0.49
37	8697.13	8697.14	-0.01	8691.66	8691.67	-0.01	4466.76	4467.12	-0.35
38	8697.18	8697.16	0.02	8830.27	8830.27	-0.00	4495.16	4495.50	-0.34
39	8878.59	8878.59	-0.00	8850.65	8850.74	-0.09	4529.49	4529.64	-0.15
40	8897.38	8897.38	-0.00	8945.48	8945.50	-0.02	4571.53	4571.66	-0.12
41	9126.09	9126.09	-0.00	8954.46	8954.47	-0.01	4623.65	4623.93	-0.28
42	9420.24	9420.24	-0.00	9009.00	9009.00	0.00	4628.67	4629.22	-0.55
43	9426.39	9426.39	-0.00	9088.01	9088.05	-0.04	4729.67	4730.04	-0.37
44	9647.10	9647.17	-0.07	9164.06	9164.08	-0.02	4734.14	4734.33	-0.19
45	9647.61	9647.61	-0.00	9214.77	9214.85	-0.08	4740.89	4741.40	-0.52
46	9745.80	9745.80	-0.00	9440.11	9440.12	-0.01	4840.27	4840.80	-0.53
47	9806.71	9806.67	0.05	9488.55	9488.54	0.01	4926.09	4926.53	-0.44
48	9806.75	9806.73	0.01	9508.91	9508.91	-0.00	4956.57	4956.91	-0.34
49	9911.10	9911.02	0.08	9619.20	9619.29	-0.09	4977.12	4978.25	-1.12
50	9911.11	9911.02	0.09	9674.67	9674.67	-0.00	4977.37	4978.44	-1.08
51	9993.68	9993.68	-0.00	9675.25	9675.25	0.00	5041.43	5042.22	-0.78
52	10004.98	10004.98	-0.00	9743.71	9743.79	-0.08	5092.00	5092.58	-0.58
53	10188.36	10188.36	-0.00	9865.30	9865.33	-0.03	5108.70	5109.39	-0.68
54	10194.51	10194.45	0.06	9922.90	9922.92	-0.02	5140.25	5141.01	-0.76
55	10208.77	10208.77	-0.00	9993.94	9993.95	-0.01	5153.10	5154.39	-1.29
56	10292.54	10292.54	-0.00	10006.55	10006.59	-0.04	5177.42	5177.65	-0.24
57	10508.45	10508.45	-0.00	10132.33	10132.38	-0.05	5186.87	5187.25	-0.38
58	10517.59	10517.59	-0.00	10165.62	10165.60	0.02	5204.40	5205.00	-0.61
59	10742.14	10742.14	-0.00	10266.38	10266.40	-0.02	5246.10	5246.39	-0.29
60	10742.73	10742.73	-0.00	10304.11	10304.10	0.01	5255.85	5256.28	-0.44
61	10842.18	10842.18	-0.00	10364.93	10364.90	0.03	5312.30	5312.86	-0.57
62	10904.69	10904.69	-0.00	10460.51	10460.50	0.01	5320.86	5322.42	-1.56
63	10904.77	10904.77	-0.00	10636.68	10636.70	-0.02	5324.95	5325.30	-0.35
64	11008.77	11008.70	0.08	10654.81	10654.87	-0.06	5357.25	5358.09	-0.84
65	11008.79	11008.77	0.01	10749.73	10749.70	0.03	5385.38	5386.43	-1.05
66	11097.17	11097.17	-0.00	10753.95	10753.98	-0.03	5415.13	5415.83	-0.70
67	11101.53	11101.53	-0.00	10758.04	10758.00	0.04	5417.34	5417.91	-0.56
68	11278.08	11278.08	-0.00	10861.98	10862.15	-0.17	5432.57	5433.31	-0.75
69	11291.96	11291.96	-0.00	10871.09	10871.10	-0.01	5462.80	5463.15	-0.35
70	11294.93	11294.93	-0.00	10922.00	10922.03	-0.03	5489.37	5490.04	-0.67
71	11391.68	11391.68	-0.00	10925.18	10925.30	-0.12	5489.92	5490.58	-0.65
72	11582.88	11582.88	-0.00	11006.52	11006.50	0.02	5532.31	5532.38	-0.06
73	11595.14	11595.14	-0.00	11035.70	11035.70	0.00	5543.86	5544.61	-0.75
74	11824.13	11824.13	-0.00	11065.18	11065.23	-0.05	5552.06	5552.88	-0.82
75	11824.65	11824.65	-0.00	11198.41	11198.51	-0.10	5625.66	5626.49	-0.83
76	11925.74	11925.75	-0.00	11225.82	11225.90	-0.08	5651.04	5651.16	-0.13
77	11990.55	11990.55	-0.00	11271.68	11271.70	-0.02	5659.48	5660.99	-1.51

78	11990.65	11990.65	-0.00	11489.39	11489.40	-0.01	5666.66	5667.46	-0.80
79	12095.40	12095.40	-0.00	11536.09	11536.10	-0.01	5680.97	5681.46	-0.50
80	12095.44	12095.44	-0.00	11539.63	11539.71	-0.08	5687.28	5688.23	-0.96
81	12149.52	12149.46	0.06	11549.66	11549.60	0.06	5718.43	5718.83	-0.40
82	12149.55	12149.52	0.04	11589.99	11590.02	-0.03	5731.64	5732.19	-0.55
83	12186.40	12186.40	-0.00	11672.67	11672.80	-0.13	5766.47	5766.91	-0.44
84	12188.55	12188.55	-0.00	11688.05	11688.10	-0.05	5768.41	5769.06	-0.65
85	12334.64	12334.64	-0.00	11710.13	11710.10	0.03	5770.76	5771.67	-0.92
86	12383.69	12383.69	-0.00	11744.32	11744.36	-0.04	5809.52	5809.85	-0.33
87	12384.60	12384.60	-0.00	11744.36	11744.50	-0.14	5822.38	5822.71	-0.33
88	12481.05	12481.05	-0.00	11832.45	11832.49	-0.04	5887.13	5888.44	-1.31
89	12524.83	12524.64	0.20	11969.83	11969.90	-0.07	5887.96	5888.55	-0.59
90	12525.35	12525.21	0.13	11970.23	11970.27	-0.04	5889.15	5890.16	-1.01
91	12643.27	12643.27	-0.00	11977.84	11977.70	0.14	5936.19	5937.06	-0.87
92	12658.92	12658.92	-0.00	12055.72	12055.70	0.02	5984.87	5985.33	-0.46
93	12695.20	12695.20	-0.00	12102.62	12102.69	-0.07	5987.32	5988.05	-0.73
94	12735.21	12735.21	0.00	12192.45	12192.50	-0.05	5996.72	5997.93	-1.21
95	12892.52	12892.52	-0.00	12200.53	12200.60	-0.07	5998.74	5999.46	-0.72
96	12892.83	12892.83	-0.00	12304.49	12304.50	-0.01	6053.12	6053.56	-0.44
97	12995.95	12995.95	-0.00	12311.73	12311.70	0.03	6091.73	6093.56	-1.83
98	13063.69	13063.69	-0.00	12374.87	12375.11	-0.24	6107.28	6108.32	-1.04
99	13063.80	13063.80	-0.00	12382.04	12382.00	0.04	6122.55	6124.52	-1.96
100	13170.37	13170.37	-0.00	12384.74	12384.75	-0.01	6177.20	6178.84	-1.63

2.7.3 Further Investigations

Convergence with respect to enhancement of normalizing-flow complexity

In spectral methods, the standard approach for improving the accuracy of computed energy levels is to increase the number of basis functions N. However, this can be computationally expensive, with memory costs scaling as N^2 and certain computational tasks, such as diagonalization of the Hamiltonian matrix, scaling up to N^3. The normalizing-flow approach offers a possibility of enhancing the expressivity of the basis functions by creating a more complex mapping, achieved by adding additional blocks. In Fig. 2.7, we show the sum of the 100 lowest vibrational energies (loss) of the HCN/HNC isomers as a function of the number of blocks in the iResNet model, for a fixed number of basis functions (at $P_{\max} = 16$).

Sensitivity of normalizing-flow coordinates to the number of target states

In this subsection, we examine the impact of optimizing normalizing-flow coordinates for different numbers of target states. In Fig. 2.8, we show the convergence of the first 200 vibrational energy levels of H_2S, computed with normalizing-flow coordinates with $P_{max} = 20$, corresponding to 506 basis functions. The coordinates were optimized to minimize the lowest 50, 100, 150, and 200 target states, respectively. As shown in Fig. 2.8, the average error of the first 200 energy levels decreases as the number of target states increases. The decrease in average energy error is accomplished by adjusting the coordinates to better converge the highest energy states in their respective subsets.

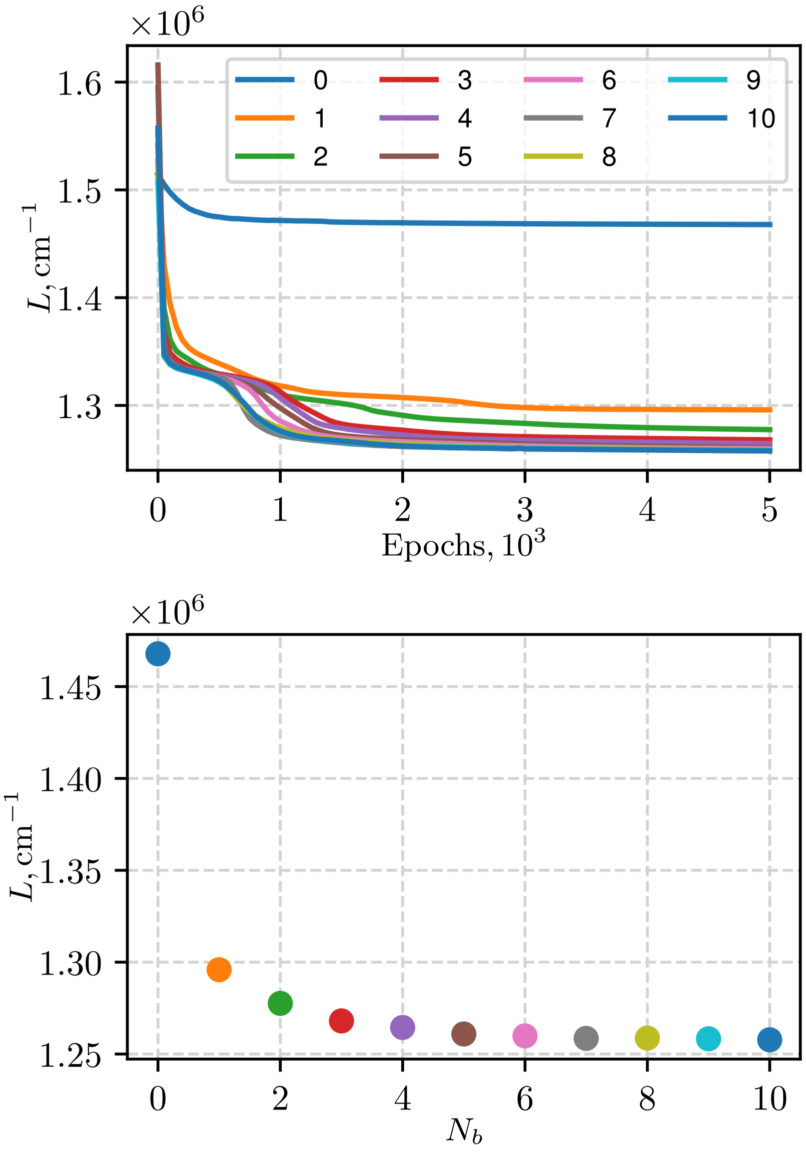

Figure 2.7: Convergence of the sum of the lowest 100 vibrational energy levels $L = \sum_i^{100} E_i$ for HCN/HNC as a function of the number of blocks in the iResNet model N_b using a basis set truncated at $P_{\max} = 16$. The upper plot shows the convergence of this sum over the number of optimization epochs, with results for different N_b represented by different colors. The lower plot shows the convergence of the optimized sum as a function of N_b.

However, as the number of target states increases, this refinement slightly compromises the convergence of the lowest energy states, which contribute progressively less to the loss function.

Transferability of the normalizing flow across different basis set truncation levels

During our investigations into the interpretability of the normalizing-flow coordinates, we noticed that the optimized coordinates remain consistent across different basis set truncations, parametrized by $P_{\max}$. This observation prompted us to explore the transferability of the learned mapping across different truncation levels. Interestingly, we found that a normalizing flow trained with few basis functions could effectively transfer to larger

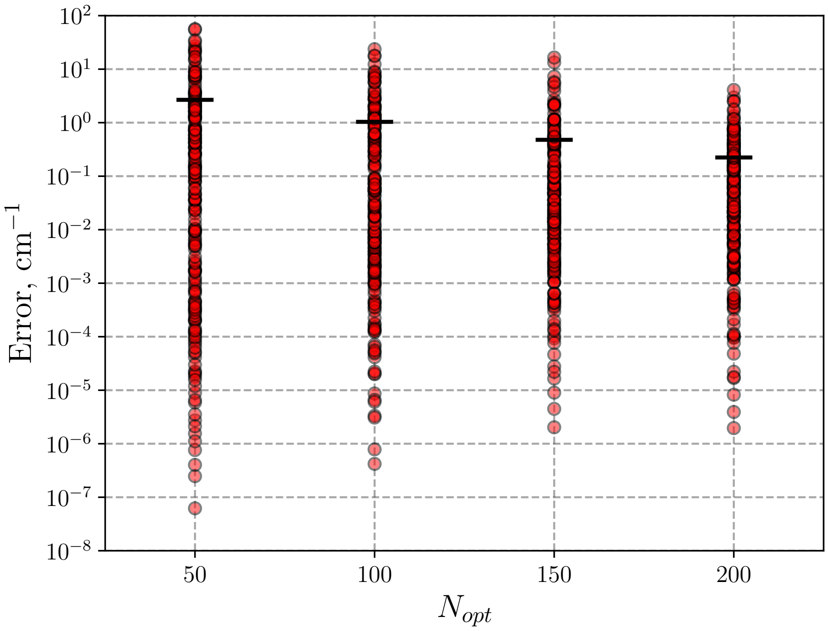

Figure 2.8: Convergence of the lowest 200 vibrational energy levels ($E_i, i = 1..200$) for H_2S, calculated using $P_{\max} = 20$ (506 basis functions) and normalizing-flow coordinates. The energy discrepancies (ΔE_i) relative to our converged benchmark reference are shown for normalizing-flow coordinates optimized for different numbers of target states (N_{opt}). The solid black lines show the average energy discrepancy.

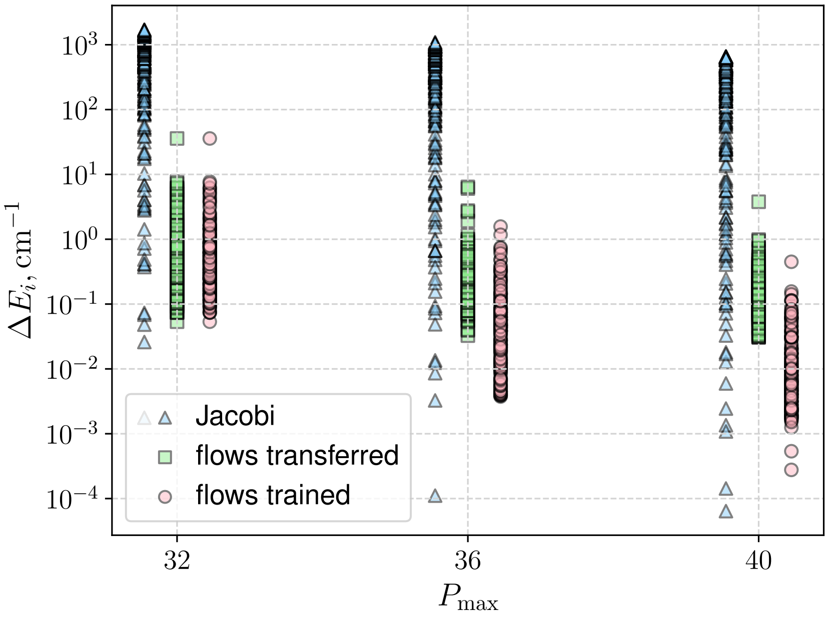

Figure 2.9: Convergence of the lowest 200 vibrational energy levels ($E_i, i = 1..200$) for HCN/HNC, calculated using Jacobi coordinates (triangles, light blue), normalizing-flow coordinates transferred from $P_{\max} = 32$ (squares, light green), and optimized normalizing-flow coordinates (circles, light pink). The energy discrepancies (ΔE_i) relative to our converged benchmark reference are shown for basis sets truncated at $P_{\max} = 32$ (1785 basis functions), 36 (2470), and 40 (3311). All states are slightly offset along the $P_{\max}$ axis for visual clarity.

basis sets without the need for retraining. In Fig. 2.9, we demonstrate this by comparing the vibrational energy levels of HCN/HNC using a normalizing flow initially trained to minimize the sum of the lowest 200 energies with a basis truncated at $P_{\max} = 32$. This mapping was then used with basis sets truncated at $P_{\max} > 32$ and compared to the results of calculations where the normalizing flow was optimized for each truncation level.

The results in Fig. 2.9 reveal that the transferred normalizing-flow mapping provides comparable accuracy to the optimized mapping for each truncation level and performs significantly better than Jacobi coordinates.

It is worth highlighting that transferability enables significant computational savings. Expanding calculations to incorporate a larger number of basis functions is feasible as long as the quantities dependent on normalizing flow, such as $\frac{\partial q_\alpha}{\partial r_l}$, $\frac{\partial D}{\partial r_l}$, $\frac{1}{D}$, etc, can be efficiently stored in memory or recomputed on the fly. Importantly, the size and computation cost of these quantities are independent of the basis size, meaning that calculations using a pre-trained normalizing flow scale with basis size in the same way as those using a regular linear mapping.

Based on these observations, we propose the following efficient protocol for implementing our approach: (i) Train a normalizing flow on a small basis set and store the optimized parameters. (ii) Transfer the learned coordinates to a larger basis set where training would be computationally prohibitive, then solve the eigenvalue problem using the pre-trained normalizing flow to obtain accurate energy levels. This protocol is particularly useful for high dimensional systems, where training costs are significantly higher. The results for H_2S and H_2CO in Table 4 where obtained using this protocol, being transferred from the optimization at $P_{max} = 12$ and $P_{max} = 9$ respectively.

Electronic calculations

The nonlinear ansatz described in Eq. 3 of the manuscript can be used to solve the electronic Schrödinger equation. To demonstrate this, we computed electronic states of prototypical one-electron systems such as the hydrogen atom, H_2^+ molecular ion, and carbon atom in the single-active electron approximation. We only consider normalizing flows for the radial coordinate and integrated out the angular coordinates using the spherical-harmonic basis. This approach is consistent with common practices in quantum chemistry, where the radial basis typically presents the primary challenge and is the main target of optimization. As a starting basis we explored Hermite basis, the basis of three-dimensional isotropic harmonic oscillator, and the standard atomic-orbital basis sets from Dunning's family [98, 99]. The number of basis functions was defined by the number of radial functions.

The Hamiltonians for the one-electron radial problems were derived by integrating out the angular coordinates using spherical-harmonic functions $Y_{l,m}$ as the angular basis, or solid-harmonic functions $r^l Y_{l,m}$ in the case of atomic-orbital basis sets. For the Hermite and isotropic harmonic oscillator radial basis sets, we utilized a direct product of the radial and angular basis functions with $l \leq l_{\max}$, where $l_{\max} = 3$ for atoms and 8 for molecules. For the atomic-orbital basis set, we employed specific combinations of radial Gaussian and solid-harmonic functions as dictated by the structure of the basis sets.

The angular matrix elements of the electron-nuclei Coulomb-attraction potential were computed using the Laplace expansion, with truncation determined by $l_{\max}$ of the angular basis.

The electronic energies of multi-electron systems, e. g., carbon atom, were calculated using the single-active-electron approximation, in which the electron-electron Coulomb-interaction potential is approximated by the one-electron electrostatic potential created by the electron density $\rho(\mathbf{r})$ of a singly-charged ion in the ground state, i. e.,

$$V(\mathbf{r}) = -\int \frac{\rho(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}'.$$

In addition, the exchange-interaction potential is generally added using some popular approximations, like the local-density approximation models. Our focus is primarily on investigating the basis set convergence of energies and its enhancement *via* normalizing flows. The absolute accuracy of the electronic energies was not our primary concern, hence we did not include the exchange-interaction potential in our calculations.

In simulations for carbon atom, the one-electron electrostatic potential of C^+ was calculated using the second-order approximate coupled-cluster CC2 level of theory with the aug-cc-pV5Z atomic-orbital basis set [98, 99], as implemented in the Psi4 quantum chemistry package [100]. The calculations were performed on a grid generated by the direct product of an equidistant radial grid $\{r_g\}_g^{N_g}$ and the Lebedev quadrature grid of 131st order represented by a set of angular points $\{\theta_h, \phi_h\}_h^{N_h}$ [101]. The radial potential was calculated by integrating angular coordinates at each radial point in the basis of spherical-harmonic functions using Lebedev quadrature rule, i. e.,

$$\begin{aligned} V_{l'm',lm}(r_g) &= \langle Y_{l'm'}(\theta,\phi)|V(r_g,\theta,\phi)|Y_{lm}(\theta,\phi)\rangle \\ &\approx \sum_h^{N_h} w_h Y^*_{l'm'}(\theta_h,\phi_h) Y_{lm}(\theta_h,\phi_h) V(r_g,\theta_h,\phi_h), \end{aligned}$$

where w_h is the Lebedev quadrature weight including spherical volume element. For computing the angular integrals between the spherical harmonics centered at different centers, which are needed in atomic-orbital-basis calculations for molecules, we also implemented Becke's partitioning scheme [102]. The calculated values of the radial potential $V_{l'm',lm}(r_g)$ were interpolated across the radial grid $\{r_g\}_g^{N_g}$ using the regular grid interpolator technique.

The reference energies for the hydrogen atom are known analytically, while for the hydrogen-molecule cation H_2^+, the energies calculated using the Riccati-Padé method [103] were employed as the reference. In Table 5, we present the absolute errors for the ground and four lowest excited electronic states of H atom and H_2^+, as calculated using normalizing flows with different basis sets. For comparison, the table also includes errors corresponding to full-configuration interaction calculations in large atomic-orbital basis sets aug-cc-pV5Z and aug-cc-pV6Z [98, 99], as computed using the Psi4 quantum chemistry package [100].

The application of the normalizing-flow ansatz with Hermite basis set demonstrates high accuracy and fast basis-set convergence for the radial electronic problem, although such basis is typically considered unsuitable for solving electronic problems. In Fig. 2.10, we show the convergence of the electronic energies for H atom and H_2^+ molecular ion as a function of the number of Hermite radial basis functions N. The results are compared against the linear parametrization, obtained by setting $g(\mathbf{x}) = \mathbf{a}\cdot\mathbf{x} + \mathbf{b}$ in Eq. 3 of the

	H atom		
State	Hermite	iso-HO	AV6Z
Ground	-4.8	-0.02	-0.0007
1st excited	-0.07	-0.002	-0.08
2d excited	-0.07	-0.002	-1.22
3d excited	-0.07	-0.002	-1.22
4th excited	-0.8	-0.002	-1.22

	H_2^+ ion	
State	Hermite	AV5Z
Ground	-3.1	-0.01
1st excited	-1.7	-0.01
2d excited	-0.007	-0.44
3d excited	-0.007	-0.44
4th excited	-0.62	-0.53

Table 5: The errors (exact − calculated, in mHartree) in the ground and excited state calculations for H atom and H_2^+ ($R_{\text{H–H}}$ = 2 Bohr), using normalizing flows with Hermite and three-dimensional isotropic harmonic oscillator (iso-HO) basis sets. For comparison, errors corresponding to full-configuration interaction approach with aug-cc-pV6Z (AV6Z) basis for H and aug-cc-pV5Z (AV5Z) basis for H_2^+ are also listed. Note that the normalizing-flow calculations of H_2^+ employed s single-center spherical-harmonic basis truncated at $l_{\text{max}} = 8$, which accounts for estimated $-1.1, -0.2$, $-10^{-5}, -10^{-5}, -0.15$ mHartree differences in the corresponding electronic energies.

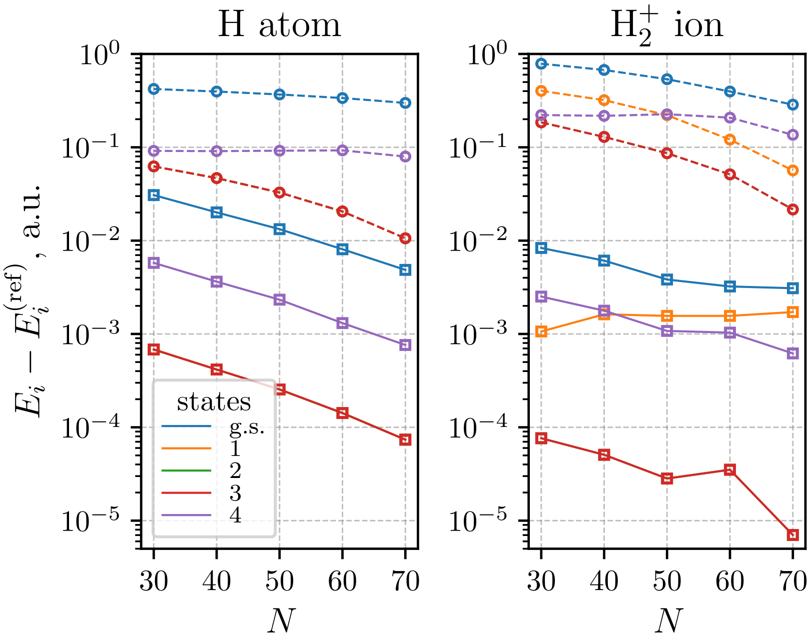

Figure 2.10: Convergence of the ground (g.s.) and four lowest excited electronic states of hydrogen atom and hydrogen molecular ion H_2^+ ($R_{\text{H–H}}$ = 2 Bohr) plotted against the number of Hermite radial basis functions N. Results obtained with optimized-linear (dashed, circles) and normalizing-flow (solid, squares) parameterizations are compared. The reference energies E_i^{ref} for H_2^+ are from Fernández and Garcia [103].

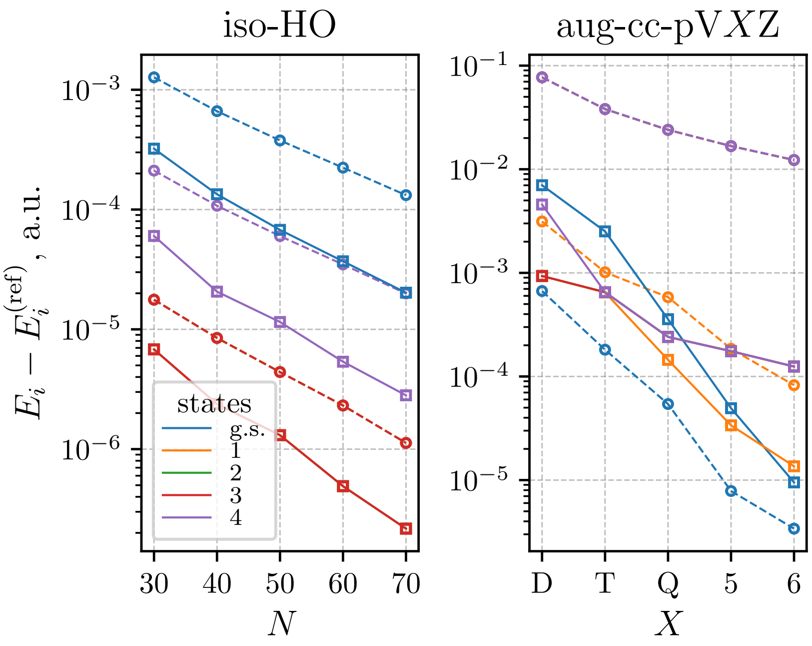

Figure 2.11: Convergence of the ground (g.s.) and four lowest excited electronic states of hydrogen atom plotted against the size of the radial basis, for the three-dimensional isotropic harmonic oscillator (iso-HO) and atomic-orbital (aug-cc-pVXZ) basis sets. Results obtained with fixed-linear (circles, dashes) and normalizing-flow (solid, squares) parameterizations are compared. Since the four excited states of the hydrogen atom are degenerate, this results in indistinguishable and overlapping error values in some of the plots.

manuscript, where the linear parameters **a** and **b** were optimized. Using the normalizing-flow ansatz, the electronic energies of H and H_2^+ converge very quickly to within few mHartree of the exact values for the ground state, respectively, and even more precisely for the excited states. The accuracy for H_2^+ is constrained to ~ 1 mHartree for the ground state and less than that for excited states due to truncation in the angular basis set. When employing a linear parametrization, a considerably larger number of Hermite-basis functions would be required to reach the same accuracy.

Normalizing flows also improved the performance of more common basis sets used in electronic structure computations, such as the basis set of isotropic three-dimensional harmonic oscillator (iso-HO) and the augmented correlation-consistent aug-cc-pVXZ (X=D, T, Q, 5, 6) atomic orbital basis sets from Dunning's family [98, 99]. Results for hydrogen atom are presented in Fig. 2.11, illustrating that the iso-HO basis set, while inherently well-suited and fast-converging for the problem, still exhibits a notable enhancement in accuracy when composed with a normalizing flow. Although the original atomic-orbital basis set demonstrates quicker ground-state convergence compared to that composed with a normalizing flow, the latter yields results that are more accurate and fast-converging for excited states. The lack of improvement in the ground-state energy by normalizing flow is likely due to the nature of atomic orbital basis sets, which consist of Gaussian functions with exponents carefully optimized for the ground state of an atom. This specific optimization likely accounts for ineffectiveness of normalizing flow in improving the ground electronic energy.

In Fig. 2.12, we show the convergence for the ground and four lowest excited electronic states of carbon atom with the number of radial basis functions, as calculated using the

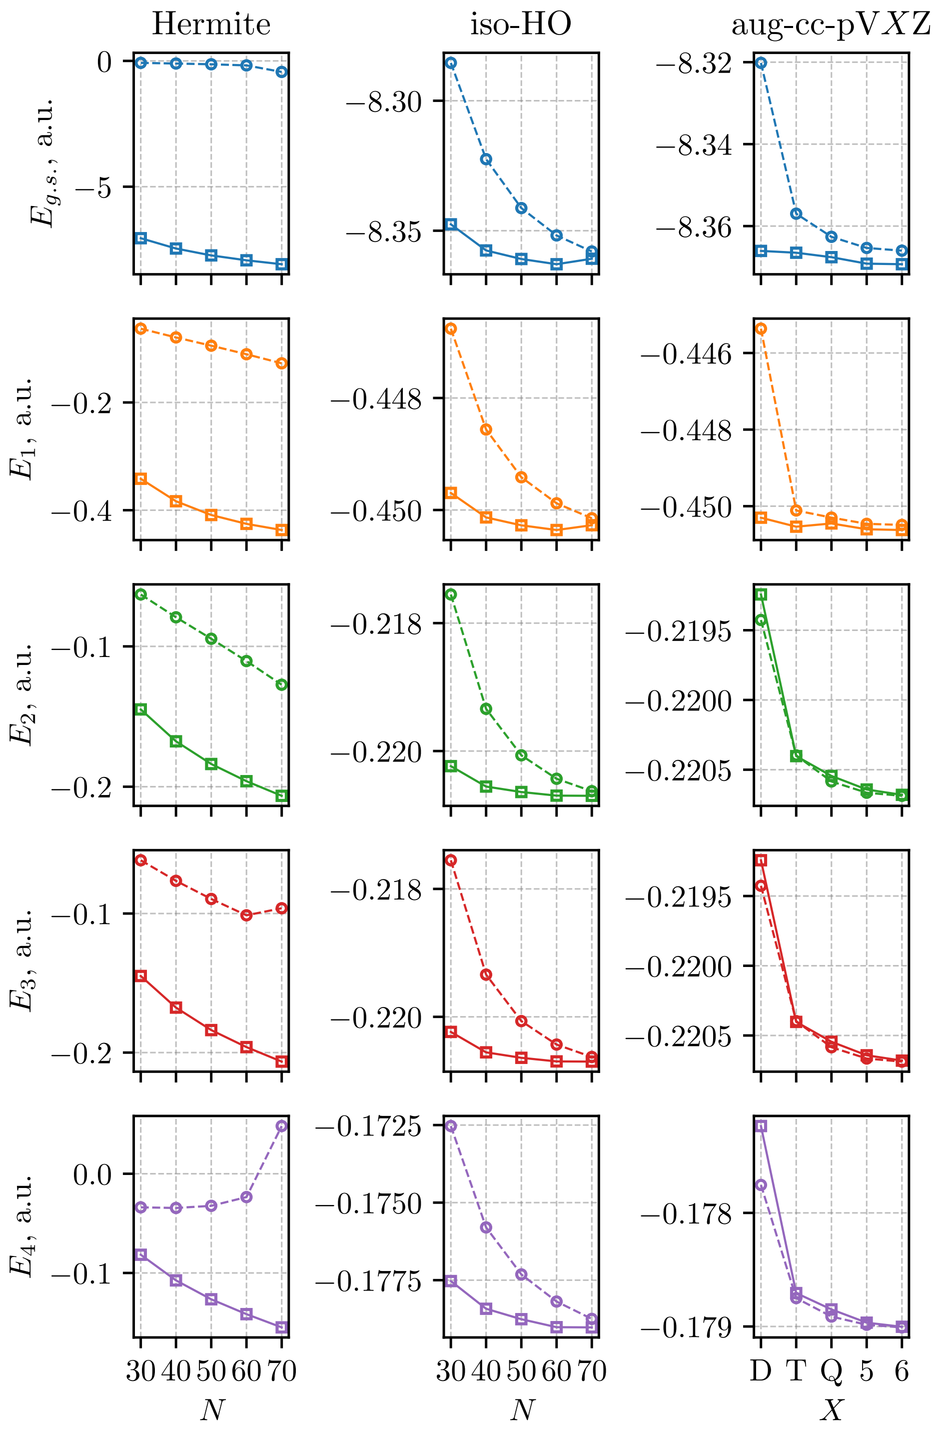

Figure 2.12: Convergence of the ground (g.s.) and four lowest excited electronic energies of carbon atom plotted against the size of the radial basis N (or X), for the Hermite, three-dimensional isotropic harmonic oscillator (iso-HO) and atomic-orbital (aug-cc-pVXZ) basis sets. Results obtained with optimized-linear (circles, dashed lines) and normalizing-flow (squares, solid lines) parameterizations are compared. The energies are calculated using the single-active-electron approximation with the Coulomb potential obtained from a quantum chemical calculation and neglecting the exchange-interaction potential.

normalizing-flow approach and an optimized-linear parametrization. The convergence is plotted for different basis sets, including the Hermite basis set, the isotropic three-dimensional harmonic oscillator basis set, and augmented correlation-consistent atomic orbital basis sets. The data clearly shows that the normalizing-flow approach generally improves the accuracy of results for all states across different basis sets, compared to optimized-linear mapping. However, an exception can be noted in the case of atomic-orbital basis sets. While normalizing-flow parametrization leads to faster convergence for the ground and first excited states, it exhibits comparative or slightly inferior accuracy for the remaining excited states, particularly with smaller basis sets, as compared to an optimized-linear mapping. We also noticed some convergence issues of the training procedure with a larger number of basis functions, which can be attributed to the limitations in the accuracy of numerical integration with Hermite and Laguerre quadratures used for calculations across all basis sets.

3 Transferability and interpretability of normalizing-flow coordinates[2]

In Chapter 2, the normalizing-flow coordinates were introduced to enhance the calculations of the solution of the TISE. We discussed the interpretation of the optimal coordinates and, based on the conclussions of the analysis, proposed the transferability between basis set truncations. This marks a starting point to obtain a deeper understanding of these coordinate transformations, while other questions were left unresolved. Why can we not achieve a one-to-one mapping between basis set functions and eigenfunctions using the normalizing-flow approach? Are there any necessary constraints on the transformations? How does the normalizing-flow coordinate change with different number of basis set functions and number of target states?

In this chapter, we adress all the questions related to the interpretability of the normalizing-flow coordinates. Notably, this analysis is not restricted to the normalizing-flow approach, but can also be applied to other coordinate transformations that are used in standard spectral methods, e. g., linear transformations.

We explore the transferability of the optimal coordinate transformations between molecules that have similar structures. This transfer is made using fine-tuning of few parameters, representing a minimal reparametrization. The success of this technique opens the possibility of the creation of a collection of pre-optimized coordinate sets that can be applied to a variety of molecules.

3.1 Introduction

Theoretical and computational studies of molecular vibrations are central to theoretical chemistry, molecular physics, and related scientific fields. Key areas of interest include the calculation and fitting of accurate potential energy surfaces (PESs) [104–107], the development of effective molecular Hamiltonians for efficient representations of molecular spectra [44, 46, 108], and the first-principles computation of accurate vibrational energies, wavefunctions, and spectra [42–44, 46]. The latter can be accomplished with perturbation theory methods, variational theory approaches using finite-basis expansions, or pseudo-variational methods like discrete-variable representations. Central to all these methods is the construction of an appropriate vibrational Hamiltonian, which relies on carefully chosen coordinates and associated basis functions. The selection of coordinates plays an important role in defining the Hamiltonian operator, influencing the extent to which vibrational motions are coupled [48, 109] and thus the overall computational efficiency. Determining an optimal coordinate system and compatible basis functions for describing molecular vibrations often requires substantial expertise and prior knowledge of the vibrational motions.

Rectilinear normal coordinates are effective for computing low-energy states in semi-rigid molecules, where the PES often can be relatively well approximated by a low-order Taylor-series expansion around a single equilibrium geometry. However, these

---

[2] This chapter is based on the manuscript by Emil Vogt, Álvaro Fernández Corral, Yahya Saleh and Andrey Yachmenev, "Transferability and interpretability of vibrational normalizing-flow coordinates". The manuscript is currently under revision in the *Journal of Computational Physics*. Additionally, it is available at arXiv `https://doi.org/10.48550/arXiv.2502.15750`. My contributions to this work were the study of transferability on finite and semi-finite domains, its practical application to the water molecule, and the collaboration to interpreting results and writing of the manuscript.

coordinates become inadequate in calculations of delocalized states, such as those encountered in floppy molecules [49, 50], or high-energy states in semi-rigid molecules [52, 53], which sample larger and more complex regions of the phase space.

Curvilinear coordinates like Radau [23], Jacobi [54, 55], valence [56], ellipsoidal [27], hyperspherical [110], and polyspherical [57–59] types are often better suited for vibrational calculations in floppy molecules [43, 60]. The optimal choice depends on the specific nuclear motions involved and requires an understanding of the morphology of the PES. Because vibrational motions vary widely across different molecular structures, no single coordinate system is universally optimal. The general approach for selecting effective coordinates is to capture the primary variations in the PES along individual dimensions, which helps to minimize coupling between different vibrational modes. Chemical intuition suggests that valence coordinates are most suitable for many molecules, as the localized electron density between adjacent atoms leads to significant changes in the PES with variations in bond lengths.

A promising strategy is to use general parametrized coordinates, with parameters optimized in variational calculations of vibrational energies. Optimized linear combinations of normal or valence coordinates have been explored in several studies [27, 61–66, 69–71]. While this strategy has demonstrated an improvement in basis-set convergence of energy calculations, it has not gained widespread adoption. The main reasons are likely that the improvements achieved through optimized linear coordinate mappings is often modest, due to their restrictive nature, and that such mappings introduce additional complexity by the loss of symmetry inherent in such coordinate systems [111].

Once the vibrational coordinates have been selected, the next step is to choose an appropriate basis set, guided by the domain of each coordinate. For example, valence coordinates have domains of $[0, \infty)$ for bond lengths, $[0, \pi]$ for angles, and $[0, 2\pi]$ for dihedrals. A common approach is to use a direct-product basis of univariate orthogonal polynomial-based functions (or linear combinations thereof), primarily due to their close relation with Gaussian quadratures [50, 112]. This approach simplifies the evaluation of integrals necessary for the calculation of the Hamiltonian matrix elements and facilitates transformations between the finite-basis representations and their corresponding discrete-variable representations. The specific univariate functions are chosen based on the coordinate domains, the shape of the PES along the coordinates, and the degree of vibrational coupling in the Hamiltonian.

We recently introduced a new approach for nonlinear parametrization of vibrational coordinates based on normalizing flows [73, 113], implemented through an invertible residual neural network [76]. In the machine learning literature, normalizing flows refer to a sequence of invertible and differentiable transformations [73, 74]. In our method, the neural network parameters are optimized using the variational principle to minimize the approximate energies. Unlike traditional approaches - where the vibrational coordinates are defined first and the basis functions are chosen afterwards - our approach adapts the coordinate system itself to maximize the performance of a chosen truncated basis set, tailoring it to the specific molecule. In our previous work [113], normalizing-flow coordinates were applied to molecules such as H_2S, H_2CO and HCN/CNH. It was demonstrated that the learned coordinates enhance the separability of the Hamiltonian, enabling more reliable assignment of approximate quantum numbers. Furthermore, we demonstrated the transferability of these coordinates across different basis-set truncations. Normalizing-flow coordinates have also been used in conjunction with Monte Carlo integration to tackle larger molecular systems, such as CH_3CN and C_2H_4O [92], and to investigate anharmonic effects in lithium solids at finite temperatures [114].

In this work, we investigate interpretability and effectiveness of normalizing-flow coordinates. Specifically, we show that when optimized using the variational principle, these coordinates shift the average density center (defined as the trace of the position operator matrix divided by the number of basis functions) to align with that of the eigenbasis. This observation offers a practical and physically meaningful way of interpreting how coordinate transformations influence the representation of quantum states. Additionally, we demonstrate that the nonlinear mappings enabled by normalizing flows offer significant advantages over linear transformations, particularly in capturing anharmonic behavior across semi-finite and finite domains.

We further investigate the transferability of normalizing-flow vibrational coordinates. Coordinates trained for H_2S are found to generalize well to its deuterated isotopologues (D_2S and HDS) with only minor, physically motivated adjustments. Remarkably, the same coordinate system also performs well for H_2O, suggesting that the learned coordinates capture structural motifs common to chemically related molecules. In all cases considered, normalizing-flow coordinates outperform traditional curvilinear valence-bond coordinates, as evidenced by faster convergence of variationally computed vibrational energies.

These findings are significant for two main reasons. First, they suggest that normalizing flows provide a practical means of transferring optimized coordinate systems across chemically related systems, thereby improving computational efficiency and initialization. Second, and more broadly, our method offers a general computational framework for identifying informative coordinate systems - those that simplify vibrational complexity and improve spectral convergence. While the notion of intrinsic coordinates is often invoked in chemical intuition to describe a minimal set of variables underlying molecular motion, our results suggest that coordinates optimized for computational efficiency may also reflect such intrinsic structures. Importantly, the optimality of these coordinates is relative: it depends on both the chosen truncated basis and the target vibrational states.

To build physical insight and isolate key effects, we focus primarily on illustrative one-dimensional systems. These simplified models serve as a clean testbed for understanding the mechanisms and interpretability of variational coordinate optimization using normalizing flows.

3.2 Theory

In variational basis representations, the vibrational Schrödinger equation,

$$\hat{H}\Psi_n = \left(\hat{T} + \hat{V}\right)\Psi_n = E_n\Psi_n, \tag{3.1}$$

is projected onto a finite set of orthonormal basis functions. Here, $\hat{H}$ is the vibrational Hamiltonian operator, $\hat{T}$ is the kinetic energy operator and $\hat{V}$ is the potential energy operator. Ψ_n and E_n are the n-th eigenfunction and eigenvalue, respectively. The vibrational eigenfunctions, Ψ_n ($n = 0 \ldots N-1$), are approximated as linear combinations of M basis functions $\{\phi\}_{m=0}^{M-1}$, with $M \geq N$, as

$$\Psi_n(\mathbf{r}) \approx \tilde{\Psi}_n(\mathbf{r}) = \sum_{m<M} c_{nm}\phi_m(\mathbf{r}), \tag{3.2}$$

where $\mathbf{r}$ denotes the vibrational coordinates.

By introducing this linear expansion into the weak formulation of the Schrödinger equation, one obtains a matrix eigenvalue problem

$$\mathbf{HC} = \mathbf{EC}, \tag{3.3}$$

where $\mathbf{H} = \{\langle\phi_m|\hat{H}|\phi'_m\rangle\}_{m,m'=0}^{M-1}$ is the Hamiltonian matrix, $\mathbf{C} = \{c_{mn}\}_{m=0,\, n=0}^{M-1,\, N-1}$ are the linear expansion coefficients, and $\mathbf{E} = \{E_n\}_{n=0}^{N-1}$ are the approximated vibrational energies. The accuracy of calculated energies can be systematically improved by increasing the number of basis functions M, ensuring variational convergence to the true energies as lower bound. In practice, quadratures or truncated Taylor-series expansions are often employed to evaluate the integrals required to construct the Hamiltonian matrix elements, which introduce additional errors to the truncated basis representation and may result in a violation of the variational principle.

An alternative approach to systematically improve calculated energies is to enhance the approximation power of the chosen basis functions. To achieve this, we start with a truncated set of orthonormal basis functions $\{\phi_m(\mathbf{q})\}_{m=0}^{M-1}$ defined on a coordinate set $\mathbf{q}$. The coordinate set $\mathbf{q}$ is related to an initial set of vibrational coordinates $\mathbf{r}$ through a parametrized invertible mapping f_θ, such that $\mathbf{q} = f_\theta(\mathbf{r})$. Because the mapping is invertible, the reverse relation also holds, $\mathbf{r} = f_\theta^{-1}(\mathbf{q})$. The augmented basis functions are then defined as

$$\gamma_m(\mathbf{q};\theta) = \phi_m(\mathbf{q})\sqrt{D}, \tag{3.4}$$

where $D = |1/\det(\nabla_{\mathbf{q}} f_\theta^{-1}(\mathbf{q}))| = |\det(\nabla_{\mathbf{r}} f_\theta(\mathbf{r}))|$ is the absolute value of the inverse of the determinant of the Jacobian. The inclusion of the factor $\sqrt{D}$ ensures that the augmented basis functions remain orthonormal, regardless of the values of the parameters θ. These basis functions can also be evaluated in the vibrational coordinates $\mathbf{r}$ as $\gamma_m(f_\theta(\mathbf{r});\theta)$.

In principle, the mapping f_θ can be any differentiable invertible function. However, for the augmented basis set to remain complete, f_θ must be bi-Lipschitz [75]. This means that there exist constants $k, K > 0$ such that $k \leq 1/|\det(\nabla_{\mathbf{q}} f_\theta^{-1}(\mathbf{q}))| \leq K$ for all $\mathbf{q}$ within the domain. We represent f_θ as a normalizing flow implemented through an invertible residual neural network (iResNet) [76], which is bi-Lipschitz by construction.

Matrix elements of the vibrational kinetic and potential energy operators can be expressed within the augmented basis in (3.4), by applying the coordinate transformation $\mathbf{q} = f_\theta(\mathbf{r})$. For the potential energy, this leads to the expression:

$$\begin{aligned} \mathbf{V}_{mm'} &= \int \phi_m^*(f_\theta(\mathbf{r}))\sqrt{D}\; V(\mathbf{r})\phi_{m'}(f_\theta(\mathbf{r}))\sqrt{D}\,\mathrm{d}\mathbf{r} \\ &= \int \phi_m^*(\mathbf{q})V(f_\theta^{-1}(\mathbf{q}))\phi_{m'}(\mathbf{q})\,\mathrm{d}\mathbf{q}, \end{aligned} \tag{3.5}$$

where the volume elements of integration are related by $\mathrm{d}\mathbf{q} = D\mathrm{d}\mathbf{r}$. This formulation shows how the change of coordinates effectively modifies the operators within matrix elements in the original basis set $\{\phi_m\}_{m=0}^{M-1}$. Therefore, optimizing f_θ for enhancing the expressivity of basis functions in (3.4) is equivalent to optimizing the coordinates in which the Hamiltonian operator is expressed, for the chosen fixed set of basis functions.

The corresponding expression for the kinetic energy matrix elements after the coordinate transformation is

$$\begin{aligned}T_{mm'} &= \int \phi^*_m(f_\theta(\mathbf{r}))\sqrt{D}\ \hat{T}(r)\ \phi_{m'}(f_\theta(\mathbf{r}))\sqrt{D}\ \mathrm{d}\mathbf{r} \\ &= \frac{\hbar^2}{2}\sum_{kl}\int \left[\left(\frac{1}{2\sqrt{D}}\frac{\partial D}{\partial q_k} + \sqrt{D}\frac{\partial}{\partial q_k}\right)\phi^*_m(\mathbf{q})\right] \\ &\sum_{\lambda\mu}\frac{\partial q_k}{\partial r_\lambda}G_{\lambda\mu}(f_\theta^{-1}(\mathbf{q}))\frac{\partial q_l}{\partial r_\mu} \\ &\left[\left(\frac{1}{2\sqrt{D}}\frac{\partial D}{\partial q_l} + \sqrt{D}\frac{\partial}{\partial q_l}\right)\phi_{m'}(\mathbf{q})\right]\mathrm{d}\mathbf{q},\end{aligned} \tag{3.6}$$

where $G_{\lambda\mu}$ are the elements of the mass-weighted metric tensor (Wilson G-matrix). In addition, the pseudo-potential term,

$$U = \frac{\hbar^2}{32}\sum_\lambda\sum_\mu \frac{G_{\lambda\mu}}{\tilde{g}^2}\frac{\partial\tilde{g}}{\partial r_\lambda}\frac{\partial\tilde{g}}{\partial r_\mu} + 4\frac{\partial}{\partial r_\lambda}\left(\frac{G_{\lambda\mu}}{\tilde{g}}\frac{\partial\tilde{g}}{\partial r_\mu}\right), \tag{3.7}$$

where $\tilde{g} = \det\left(\mathbf{G}^{-1}\right)$, also contributes to the exact kinetic energy operator. The pseudo-potential arises from the original coordinate transformation from Cartesian to the initial coordinates $\mathbf{r}$. As the pseudo-potential is a scalar function of the vibrational coordinates, its associated matrix elements can be evaluated analogously to the potential energy operator in (3.5).

3.2.1 Invertible neural networks

To model the normalizing flow f_θ, we used an iResNet consisting of 10 blocks (five blocks for the one-dimensional examples, *vide infra*). An iResNet is given by concatenating blocks of the form

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{h}_k(\mathbf{x}_k;\theta), \tag{3.8}$$

where $\mathbf{x}_k$ is the input to the block and $\mathbf{h}_k(\mathbf{x};\theta)$ is a feed-forward neural network composed of weights, biases and nonlinear activation functions. Each block was constructed as a dense neural network with two hidden layers with unit sizes $[8, 8]$, and an output layer with the number of units equal to the number of coordinates. A block is guaranteed to be invertible if it has a Lipschitz constant < 1. The inverse of each block was obtained by fixed-point iterations.

To guarantee that the feed-forward networks $\mathbf{h}_k$ for $k = 0, \ldots, K-1$ satisfy the aforementioned Lipschitz condition, we used the LipSwish activation function

$$\sigma(x) := \frac{x}{1.1}\frac{1}{1+\exp(-x)}, \tag{3.9}$$

which has a Lipschitz constant of ~ 1. With this activation function, the block $\mathbf{h}_k$ is guaranteed to be Lipschitz if each of its weight matrices, $\mathbf{W}$, has a spectral norm < 1. This is achieved by setting

$$\mathbf{W} = \begin{cases}\mathbf{W} & \text{if } \|\mathbf{W}\|_2 < c \\ \mathbf{U}\tilde{\mathbf{\Sigma}}\mathbf{V}^T & \text{if } \|\mathbf{W}\|_2 \geq c,\end{cases} \tag{3.10}$$

where $0 < c < 1$ is a hyperparameter, $||\mathbf{W}||_2$ is the spectral norm of the weight matrix $\mathbf{W}$, $\mathbf{U}$ and $\mathbf{V}$ are the left and right singular vectors of $\mathbf{W}$, respectively, and $\tilde{\mathbf{\Sigma}}$ is a diagonal matrix containing the corresponding modified singular values,

$$\tilde{\Sigma}_{ii} = \begin{cases} \Sigma_{ii} & \text{if } \Sigma_{ii} < c \\ c & \text{if } \Sigma_{ii} \geq c. \end{cases} \tag{3.11}$$

In practice, we set the upper bound on c to be 0.9 for numerical stability.

A hyperbolic tangent wrapper was applied to the output of the inverse pass of the iResNet (input of the forward pass), $\mathbf{x}_K \longrightarrow \tanh(\mathbf{x}_K)$, which maps all dimensions to a domain of $[-1, 1]$. The final output of the inverse pass was linearly scaled from the range $[-1, 1]$ to match the domain of the original vibrational coordinates:

$$f_\theta^{-1}(\mathbf{x}) = \mathbf{a} \cdot \tilde{f}^{-1}(\mathbf{x}) + \mathbf{b}. \tag{3.12}$$

This ensures that $f_\theta^{-1}(\mathbf{x})$ is contained within the original coordinate domain for all possible values of the parameters θ.

3.2.2 Loss function

To optimize the parameters θ of f_θ, we take advantage of the variational principle to define the loss function as

$$\mathcal{L}_\theta^M = \sum_{n<M} E_n \longrightarrow \min_\theta . \tag{3.13}$$

If the number of target states is equal to the number of basis functions, $M = N$, this loss function can be efficiently computed as $\mathcal{L}_\theta^N = \mathrm{Tr}(\mathbf{H})$, eliminating the need to calculate off-diagonal matrix elements of the Hamiltonian matrix. Despite the added complexity, the computational cost of calculating off-diagonal matrix elements and repeated Hamiltonian matrix diagonalization can be justified when the optimization focuses on a specific subset of states of interest. In such cases, focusing on a smaller set of eigenvalues and selectively improving their accuracy can lead to a more efficient and targeted optimization process.

The loss function was optimized using the `Optax` [115] Adam optimizer with a learning rate of 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$, and $\bar{\epsilon} = 0.0$.

3.2.3 Details of multidimensional calculations

For the multidimensional calculations on H_2S and H_2O, the reference vibrational coordinates were chosen as conventional displacement-based valence coordinates, i. e., bond lengths and angles. Multidimensional basis functions were expressed as direct products of Hermite basis functions in all examples. The normalizing-flows architecture enables mapping of any input coordinate range, defined by the domain of $\mathbf{r}$ (initial valence coordinates), to any output coordinate range, defined by the domain of $\mathbf{q}$ (optimized coordinates). Therefore, Hermite basis functions are suitable for both the bond stretching coordinates, $(-\infty, \infty) \longrightarrow [0, \infty)$, and the angular coordinates, $(-\infty, \infty) \longrightarrow [0, \pi]$.

The direct-product basis was truncated by including only basis-product configurations (n_1, n_2, n_3) that satisfy the polyad condition $2n_1 + 2n_2 + n_3 \leq P_\text{max}$, where n_1, n_2, and n_3 represent the Hermite basis function indices corresponding to the two stretching and one bending valence coordinate, respectively. Two direct-product quadrature grids were used in an alternating fashion during training, with 30/31 points along the stretching coordinates and 60/61 points along the bending coordinate. These grids were chosen to minimize numerical integration error in the Hamiltonian matrix elements, keeping the optimization stable while maintaining computational efficiency. For higher-dimensional systems, it may be necessary to use sparse quadrature methods or Monte Carlo integration.

The PES used for H_2S and its isotopologues was obtained from Azzam et al. [83], while the PES for H_2O was taken from Conway et al. [116].

3.3 Results & Discussion

3.3.1 Solving the Morse oscillator with harmonic oscillator basis functions

To better understand the morphology and effectiveness of the normalizing-flows coordinate optimization, we examine its application to a typical one-dimensional example of molecular vibrations: the Morse oscillator [117]. The Hamiltonian for the Morse oscillator is

$$H_M(r) = \frac{-\hbar^2}{2\mu}\frac{\partial^2}{\partial r^2} + D_e\left[1 - \exp(-a_M(r - r_e))\right]^2, \tag{3.14}$$

where μ is the reduced mass of the oscillator, D_e is the dissociation energy, a_M is a second Morse parameter, and $r - r_e$ is the displacement coordinate relative to the equilibrium bond distance, r_e. The nonlinear coordinate transformation $z = 1 - \exp(-a_M(r - r_e))$ maps the Morse potential into a harmonic potential,

$$V_M(r) = D_e\left[1 - \exp(-a_M(r - r_e))\right]^2 \rightarrow \tilde{V}(z) = D_e z^2. \tag{3.15}$$

It may seem that the Hermite functions - eigenfunctions of the quantum harmonic oscillator - could be used to exactly solve the problem. However, this coordinate transformation also modifies the kinetic energy operator. Specifically, the change in the volume element $\mathrm{d}r \rightarrow \mathrm{d}z$ results in the following kinetic energy operator expressed in the z coordinate

$$\tilde{T}(z) = \frac{-\hbar^2}{2\mu}\left(a_M^2(1 - z)^2\frac{\partial^2}{\partial z^2} - a_M^2(1 - z)\frac{\partial}{\partial z}\right). \tag{3.16}$$

This expression differs from the standard kinetic energy operator for both the Morse and harmonic oscillators. It is therefore not possible to map the Morse-oscillator Hamiltonian into the harmonic-oscillator Hamiltonian by a change of coordinates. This limitation can also be realized from the fact that a coordinate transformation does not change relative positions of nodes between different basis functions. The odd harmonic-oscillator eigenfunctions share a node at the center of their expansion, but the Morse oscillator eigenfunctions do not. In addition, the transformed coordinate z is defined in the domain $(-\infty, 1)$, which is incompatible with the domain $(-\infty, \infty)$ required for harmonic-oscillator eigenfunctions.

The optimal coordinate transformation can be determined by minimizing the loss function in (3.13). As demonstrated, this transformation does not establish a one-to-one correspondence between the Hermite basis functions and the Morse-oscillator eigenfunctions. Instead, the optimal transformation ensures that the span of the N Morse-oscillator eigenfunctions of interest, $\{\Psi(r)\}_{n=0}^{N-1}$, is optimally captured within the span of the M augmented harmonic-oscillator basis functions, $\left\{\gamma\left(f_\theta(r);\theta\right)\right\}_{m=0}^{M-1}$.

We optimized linear and nonlinear coordinate transformations to obtain variational solutions of the Morse-oscillator problem. The potential parameters used were $a_M = 2.1440$ Å^{-1} (1 Å $= 10^{-10}$ m), $D_e = 42301$ cm^{-1}, representative of a typical OH-stretching mode in molecules [118]. Both coordinate transformations were optimized to approximate all 23 bound states using a truncated set of 23 Hermite functions. The normalizing-flows transformation was modelled by the iResNet with five residual blocks, which proved sufficient to achieve convergence of the coordinate.

The Morse potential expressed in fixed linear (*vide infra*), optimized linear, and normalizing-flow coordinates are shown in Fig. 3.1. In both of the optimized coordinates, the Morse potential is shifted relative to the Hermite basis reference at $r - r_e = 0$, positioning the basis center closer to the average density center of the eigenstates, $\frac{1}{N}\sum_n \langle \Psi_n | r | \Psi_n \rangle$. This shift reflects that the optimal mapping is not a direct one-to-one correspondence between the Hermite basis functions and the Morse-oscillator eigenfunctions, $\gamma_n \rightarrow \Psi_n$. As a result, the largest contribution to the n-th eigenstate does not necessarily come from the n-th basis function. The largest difference between the linear and normalizing-flow coordinates is seen at large values of r (approximately for $r - r_e > 1$ Å). In this region, the harmonic potential and the Morse potential differ significantly, leading to different decay behaviour of the respective eigenfunctions.

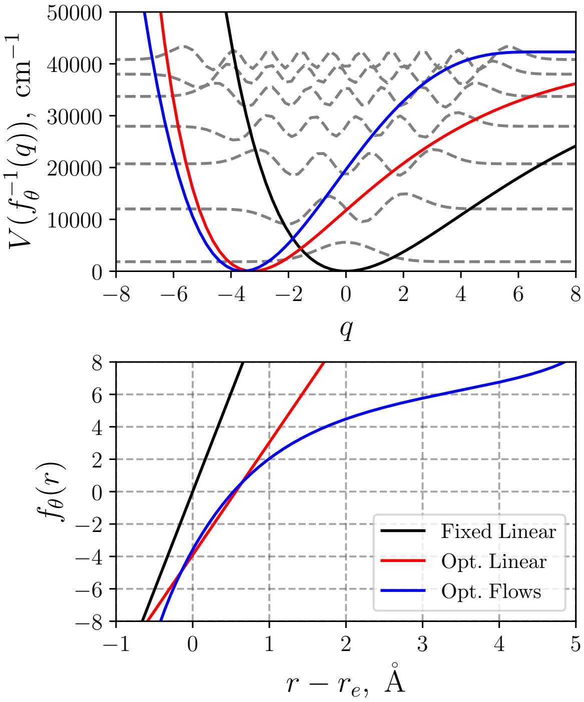

Figure 3.1:
(Top panel) Morse potential expressed in fixed linear coordinates (black), optimized linear coordinates (red), and normalizing-flow coordinates (blue). The Hermite functions, $\mathcal{H}_n(q)$, with $n = 0, 3, \ldots, 18$ are depicted as grey dashed lines.

(Bottom panel) The corresponding optimized coordinate $q = f(r;\theta)$ as a function of the displacement coordinate $r - r_e$.

The computed energy values for the Morse-oscillator problem expressed in fixed linear, optimized linear, and normalizing-flow coordinates are presented in Table 6. The parameters a and b for the fixed linear coordinates are derived from a quadratic expansion of the potential in r centered around r_e, followed by mapping of the resulting Hamiltonian onto the harmonic oscillator Hamiltonian in q. This yields,

$$a = \left(\frac{F}{\mu\hbar^2}\right)^{1/4}, \quad b = r_e, \tag{3.17}$$

where $F = \frac{d^2V(r)}{dr^2}|_{r=r_e}$ is the force constant. The fixed linear transformation is effective for the lowest energy states, where wavefunctions remain largely confined within the region where the quadratic expansion of the potential holds. However, its accuracy declines rapidly with increasing level of excitation. In contrast, optimized linear coordinates perform worse for the lowest energy states but degrade more gradually in accuracy with increasing level of excitation. Normalizing-flow coordinates offer significantly improved accuracy across all energy levels, demonstrating the advantages of nonlinear coordinate transformations.

State	Fixed linear	Opt. linear	Opt. flows	Reference
0	0.00	0.09	0.00	1838.97
1	0.00	0.70	0.00	5394.32
2	0.00	6.16	0.00	8786.20
3	0.00	34.96	0.00	12014.61
4	0.01	99.80	0.00	15079.56
5	0.48	174.38	0.00	17981.04
6	9.76	263.49	0.00	20719.06
7	95.06	439.24	0.00	23293.61
8	458.58	712.07	0.00	25704.69
9	1312.18	1013.82	0.00	27952.31
10	2728.98	1279.12	0.01	30036.46
11	4729.23	1472.23	0.04	31957.14
12	7279.19	1583.27	0.11	33714.36
13	10473.30	1592.80	0.27	35308.11
14	14373.51	1651.39	0.53	36738.40
15	18484.27	1609.12	0.92	38005.22
16	24442.18	2375.75	1.47	39108.57
17	36614.07	3028.89	2.14	40048.46
18	58273.63	4412.65	2.78	40824.88
19	92645.97	6551.87	3.15	41437.83
20	145985.88	7988.28	3.65	41887.32
21	231296.45	13107.43	4.83	42173.35
22	382860.60	14762.55	11.53	42295.90
Sum	1032063.33	64160.06	31.43	652300.37

Table 6: Calculated bound state energy level discrepancies (in cm^{-1}) for the Morse potential using a basis of 23 Hermite functions. Results are provided for fixed linear, optimized linear, and normalizing-flows coordinate transformations. The reference energy levels were calculated analytically.

The linear parameters in (3.17) provide a well-established approach for modifying the coordinate mapping in response to changes in physical parameters of the system (μ, F, and r_e). However, the corresponding mapping strategy for the normalizing-flow coordinates is not entirely clear. To address this, in subsection 3.3.5 we investigate the transferability of the optimized normalizing-flow coordinates in isotopologues, i. e., molecules that share the same structure but differ in nuclear masses.

3.3.2 Transformations on finite and semi-finite domains: The need for nonlinear maps

As illustrated in the previous subsection, the optimal coordinate ensures that the average density center of the basis functions is closely aligned with the average density center of the eigenfunctions. On infinite domains, this transformation can be conveniently accomplished using a simple linear transformation. However, on semi-finite and finite coordinate domains, linear transformations may produce unphysical results.

Consider a one-dimensional Schrödinger equation defined on a domain $[r_{\min}, r_{\max}]$, where at least one of the limits is finite. This constrain commonly arises in radial and angular coordinates. Let $\{\phi_i\}_{i=0}^{\infty}$ be a basis set that is optimal for a given system on this domain, e. g., the eigenfunctions of the system. Suppose we wish to reuse this basis set for a different system defined on the same interval $[r_{\min}, r_{\max}]$, but with a shifted potential minimum and a different potential width. In principle, it is possible to account for such changes using a linear transformation of the form $L(r) = ar + b$, as suggested in (3.17). This transformation maps the original domain to a new domain, $L : [r_{\min}, r_{\max}] \to [q_{\min}, q_{\max}] = [L(r_{\min}), L(r_{\max})]$. Depending on the values of a and b, the transformed domain may extend beyond the original physical region, i.e., $q_{\min} < r_{\min}$ and/or $q_{\max} > r_{\max}$, resulting in an unphysical domain. Even if the domain remains within bounds, the transformation may alter the implicit (Dirichlet or Neumann) boundary conditions of the basis functions, compromising their validity.

Linear transformations on the infinite real line preserve the domain and avoid the aformentioned challenges on semi-finite or finite intervals. We exploit this property to construct coordinate maps between two different problems defined on the same (semi-)finite domain. We define a map $h : [r_{\min}, r_{\max}] \to [r_{\min}, r_{\max}], \quad r \mapsto q = h(r)$ as a composition of invertible transformations:

$$h = L_0^{-1} \circ T^{-1} \circ L \circ T \circ L_0, \tag{3.18}$$

where $h = f \circ g$ denotes function composition $h(x) = f(g(x))$. The fixed linear transformation L_0 maps the initial domain $[r_{\min}, r_{\max}]$ to the interval I, where I is the domain of the nonlinear transformation $T : I \to (-\infty, \infty)$. After these two operations, the problem has been mapped into an infinite domain. Hence, any linear transformation L preserves the domain. The inverse transformation, T^{-1} followed by L_0^{-1}, returns the coordinate to the original physical domain $[r_{\min}, r_{\max}]$. This construction defines a nonlinear map that contains an optimizable unconstrained linear transformation L, while remaining constrained to the physical domain. We use the construction in (3.18) in subsection 3.3.6 as a fine-tuning mechanism for transfering coordinates from H_2S to H_2O.

To illustrate the limitations of linear coordinate transformations and the utility of the nonlinear transformation h, we consider a one-dimensional example defined on a finite interval. We construct two one-dimensional model potentials by fixing the valence radial coordinates at their equilibrium values and treating the valence angular coordinate as the variable. The systems considered are H_2S and H_2O, with the valence angular coordinate defined on the domain $[0, \pi]$. We treat the angular coordinate for H_2S as the one optimized for a chosen basis set and aim to reuse this basis and coordinate definition for H_2O. Rather than transforming the basis functions directly, which is difficult to visualize due to their infinite extent, we instead transform the coordinate in which the potential is expressed. This approach is analogous to transforming Hermite functions so that the second-order Taylor series of the potential at the minimum maps to that of the harmonic oscillator, as in (3.17). The two potentials are illustrated in the top panel of Fig. 3.2, which reveals differences in both equilibrium positions and potential shapes.

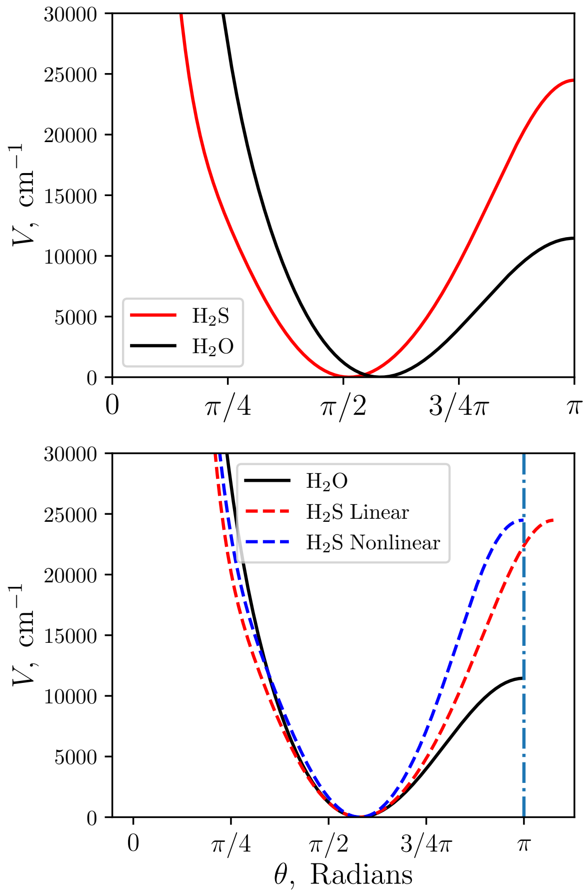

Figure 3.2:
(Top panel) The one-dimensional bending PES of H_2S (red) and H_2O (black) with bond lengths fixed to the equilibrium values.

(Bottom panel) A comparison between the H_2O PES (black) and the composition of the H_2S with two different transformation optimised to try to approximate the water potential. The linear transformation L^{-1} is shown in red and the nonlinear h^{-1} in blue. The linear approach extends the potential beyond the physical domain, marked by the light-blue vertical line at $\theta = \pi$, while the nonlinear transformation keeps the potential within the domain.

We apply two different invertible transformations to the H_2S potential, each optimized to approximate the H_2O potential. That is, we seek

$$V_{H_2S} \circ g^{-1} \approx V_{H_2O}, \tag{3.19}$$

where g is either a linear transformation L or the nonlinear transformation h defined in (3.18). Both transformations are parametrized by only two adjustable parameters. The wrapper T is chosen as the inverse hyperbolic tangent function.

The results of the optimizations are presented in the bottom panel of Fig. 3.2 alongside the original H_2O potential as a function of the valence-bond angle, which serves as the objective function. Both L and h succesfully shift the minimum and adjust the width of the composed H_2S potentials to match that of H_2O. Although the linear transformation produces a closer visual match to the H_2O potential, it also maps part of the domain outside the physically valid interval due to the absence of boundary constraints on the linear parameters. Applying such a transformation in vibrational calculations for H_2O would produce accurate results for only few low-energy states. Higher-energy eigenstates would sample unphysical regions of the Hamiltonian, resulting in unreliable or incorrect predictions.

This problem is analogous to known challenges in using Hermite basis functions for semi-finite domains, e. g., bond-stretching vibrations. While in practice, numerical stability is often retained due to the steep rise of the potential at short bond length, this is not guaranteed. In contrast, our use of wrapping functions, i. e., invertible functions that map finite intervals to infinite ones such as T, within the normalizing-flow framework allows for unrestricted coordinate transformations on infinite, semi-infinite, and finite intervals while remaining withing the valid domain.

Although this example illustrates the limitations of direct linear coordinate transformations, it neglects the effects of coordinate transformations on the kinetic energy operator. Since our goal is to minimize computed approximate energies, any coordinate transformation must account for both potential and kinetic contributions. In some applications, such as isotopologue calculations explored in subsection 3.3.5, the systems share the same PES but differ in their kinetic operators. In such cases, a transformation of the form (3.18) remains suitable for mapping between solutions.

3.3.3 Intuition of the shape of normalizing-flow coordinates

In this section, we solve the one-dimensional Schrödinger equation for a simplified double-well potential, commonly used to model the inversion vibration in ammonia (NH_3). The Hamiltonian is defined as follows:

$$H = -\frac{\hbar^2}{2\mu}\frac{\partial^2}{\partial r^2} + \frac{k(r^2 - r_0^2)^2}{8r_0^2}, \tag{3.20}$$

where $\mu = \frac{3m_H m_N}{3m_H + m_N}$, $k = \frac{8V_0}{r_0^2}$, $r_0 = 0.3816$ Å and $V_0 = 2028.6\ \mathrm{cm}^{-1}$ is the inversion barrier.

We solve the one-dimensional Schrödinger equation using the ansatz introduced in (3.4), optimizing the invertible transformation to minimize the energies of the first few states. We use as many basis set functions as there are target states and repeat the optimization for different numbers of states.

In the top panel of Fig. 3.3, we plot the optimized normalising-flow coordinate transformations $q = f(r;\theta)$ obtained from these different optimizations. When optimizing for only the two lowest states, the transformation develops a plateau around $r = 0$, where the gradient tends toward zero, and therefore so does the Jacobian determinant D. In the normalizing flow coordinate q, the plateau is centered around $q|_{r=0} = f_\theta(0) = 0$, so the composed basis function $\gamma_0(q;\theta) = \phi_0(q)\sqrt{D}$, shown in the middle panel of Fig. 3.3, becomes small near the origin. This occurs despite the primitive basis function $\phi_0(q) \propto \exp(-q^2/2)$ having a maximum at $q = 0$, and leads to a local minimum in $\gamma_0(q;\theta)$ at this point. As r moves away from zero, the factor $\sqrt{D}$ increases faster than the exponential decay of the primitive basis function $\exp(-q^2/2)$, until the latter eventually dominates. This interplay between the gradient of the coordinate transformation and the shape of the primitive basis function creates a bimodal structure in the first composed basis function, starting from an originally unimodal Gaussian function. The first two eigenfunctions of the double-well potential are symmetric and antisymmetric in r, with near-zero and zero density at $r = 0$, respectively. This explains why the coordinate transformation creates a plateau in this region and why this improves the approximation power of the transformed basis compared to the original one.

As the number of target states increases, the mean density at $r = 0$ across all eigenstates also increases (see the bottom panel of Fig. 3.3). If the same transformation optimized for just the lowest two states were retained, all basis functions would have reduced amplitude around $r = 0$, making it difficult to accurately represent higher-energy eigenstates, particularly those above the inversion barrier. For this reason, the optimized transformation becomes more linear, and the first composed basis function increasingly resembles a simple Gaussian.

Additionally, the top panel of Fig. 3.4 shows that the variation of the optimized coordinate transformation diminishes as the number of basis functions increases. For sufficiently large basis sets, the transformation exhibits minimal variation and stabilizes. This behaviour reflects the transferability of normalizing-flow coordinates across different basis-set truncations and helps to explain the ability of the method to consistently produce accurate vibrational energies.

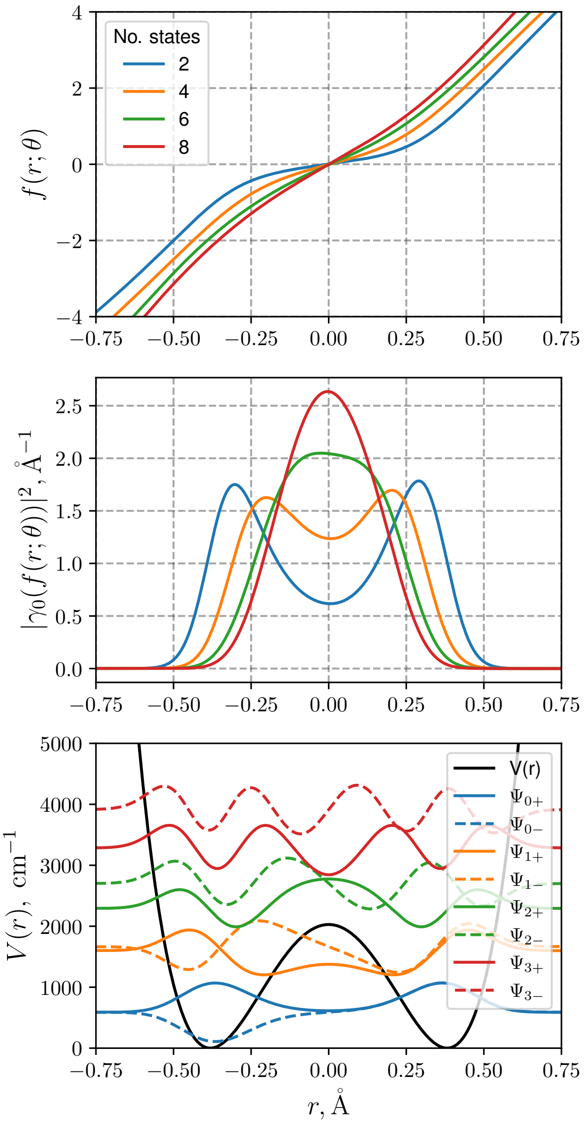

Figure 3.3:
(Top panel) Optimized normalizing-flow coordinates for varying number of target states (using the same number of basis functions as target states in each case) as a function of the physical coordinate r.

(Middle panel) The first augmented density distribution as a function of r.

(Bottom panel) Potential energy curve along with the first few first associated eigenfunctions used in the coordinate optimization. Each eigenfunction is vertically shifted by its corresponding eigenvalue for visual purposes.

3.3.4 Transferability across basis-set truncations

During the training of the normalizing-flow mapping, the neural network leverages all available information from the truncated basis set to enhance its performance. As a consequence, the training disregards basis functions that are not included in the approximation, as exemplified in the previous subsection. This might suggest that the resulting optimal coordinate mapping depend strongly on the number of basis functions employed during training. However, due to the inherent properties of normalizing flows, the transformed basis set retains completeness and orthogonality. As a result, the optimal coordinate mapping remains largely independent of the number of basis functions M, provided the basis set used for training is sufficiently expressive.

The computational cost of training scales with the number of basis functions, M, as M^3, if the loss function requires solving an eigenvalue problem using full diagonalization techniques. It also depends on the complexity of the neural network. The approximate independence from M enables an efficient training protocol, introduced in Ref. 113, in which the normalizing-flow coordinates are first optimized using a small or moderate number of basis functions. These optimized coordinates can then be transferred to calculations with larger basis sets for higher accuracy in energy predictions. We refer the reader to Ref. 113 for more details on this transferability approach. In subsections 3.3.5 and 3.3.6, we make use of this property to reduce computational cost while maintaining accuracy.

3.3.5 Transferability across isotopologues

In this section, we investigate the effectiveness of transferring normalizing-flow coordinates between isotopologues of H_2S. The coordinates were initially optimized in variational calculations for the lowest 100 vibrational energy levels of H_2S using a basis set truncated at $P_{\text{max}} = 12$, corresponding to 140 basis functions. The optimized coordinates were then applied in calculations of the vibrational energy levels of D_2S and HDS. Non-Born-Oppenheimer effects are neglected, and the PES is therefore approximated to be invariant across isotopologues.

In Fig. 3.4, we show a direct comparison of vibrational energies computed using the transferred normalizing-flow coordinates with those obtained using a fixed linear transformation of valence coordinates. In contrast to the transferred normalizing-flow coordinates, the fixed linear transformation inherently accounts for mass changes when applied to different isotopologues. Overall, the transferred normalizing-flow coordinates yield more accurate energy levels compared to the valence coordinates, with error reductions of up to three orders of magnitude. Additionally, they show faster convergence with respect to the basis-set truncation P_{max}. This improvement in accuracy is more pronounced for HDS, particularly for lower vibrational levels, which are generally better converged. We attribute this improved accuracy to the enhanced separability of the Hamiltonian when expressed in normalizing-flow coordinates [113].

Enhanced separability means that the vibrational energy levels can be more accurately approximated as sums of contributions from individual vibrational modes. However, isotopic substitution affects this separability. For example, when transitioning from H_2S to HDS, the Hamiltonian separability is slightly reduced. In such cases, the transferred flow coordinates can still provide reasonably accurate predictions for low-energy states, although accuracy diminishes for higher vibrational excitations. When both hydrogen atoms are replaced, as in the transition from H_2S to D_2S, the Hamiltonian becomes even less separable. As a result, the transferred coordinates are less effective, and the accuracy of energy predictions declines, even for the low-energy vibrational states.

To further improve the performance of the transferred flow coordinates, we fine-tuned the linear model parameters $\mathbf{a}$ and $\mathbf{b}$ in (3.12). Due to their greater mass, the wavefunctions of the heavier isotopologues are more localized, thus allowing optimization of the linear parameters without risk of mapping outside the physical domain. This

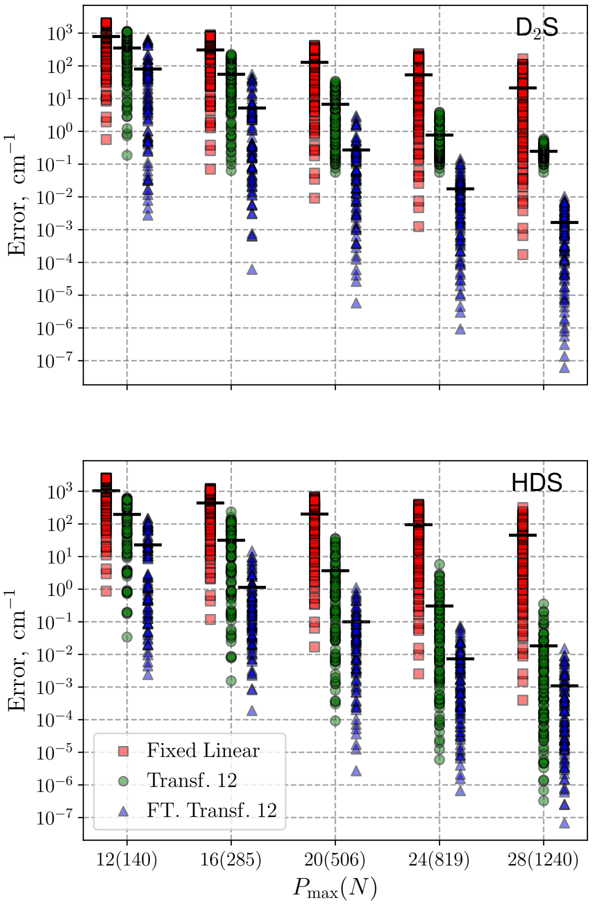

Figure 3.4:
Convergence of the first 100 vibrational energy levels of D_2S and HDS using fixed linear transformation of valence coordinates (red squares), transferred normalizing-flow coordinates (green circles), and fine-tuned transferred normalizing-flow coordinates (blue triangles). The vertical axis represents the calculated energy error, $E_i - E_i^{(\mathrm{Ref})}$, while the horizontal axis shows the basis-set truncations, P_{max}. The transferred coordinates were originally optimized for H_2S with $P_{\mathrm{max}} = 12$, and the fine-tuned parameters were obtained through re-optimization of the linear parameters in the normalizing-flows model. The black horizontal lines show the average error per energy level. The data points are slightly offset along the P_{max} axis for visual clarity.

eliminates the need for the more complex nonlinear optimization described in subsection 3.3.2. For D_2S and HDS, the fine-tuning parameters we obtained by solving the Schrödinger equation for 100 vibrational states, using just 40 iterations with the Adam optimizer. The fine-tuning was performed only for the smallest truncated basis set, $P_{\mathrm{max}} = 12$. The resulting optimized coordinates were subsequently transferred to calculations with larger truncated basis sets, without reoptimization of the parameters, as described in subsection 3.3.4. The results for the fine-tuned coordinates in Fig. 3.4 show a significant improvement of computed energy levels, by up to more than three orders of magnitude, compared to the results obtained using a fixed linear transformation of valence coordinates.

The rationale for fine-tuning transferred coordinates across isotopologues is well founded, as the fixed linear mapping of valence coordinates already includes mass-dependent scaling (see (3.17)). Based on this, the variation in the linear parameters of the normalizing-flow model is expected to scale proportionally to the diagonal elements of the mass-weighted metric tensor, i. e., $a_i \propto G_{ii}^{1/4}$ and $b_i \rightarrow 0$ as $G_{ii} \rightarrow 0$, consistent with the eigenstate average density center argument discussed in subsection 3.3.1. For valence-stretching vibrations, $G_{ii} = \mu^{-1}$, as used in (3.17).

3.3.6 Transferability across molecules

Conventional geometrically defined curvilinear coordinates are broadly applicable across various molecules. For example, valence coordinates are particularly well suited for semi-rigid molecules like H_2S, H_2O and H_2CO, and can be conveniently expressed using a Z-matrix representation. In the previous subsection, we demonstrated how optimized normalizing-flow coordinates can be effectively transferred across isotopologues. Here, we extend this analysis to evaluate the performance of normalizing-flows coordinates when transferred between different molecules with similar chemical structures.

In Fig. 3.5, we compare the convergence of the first 100 vibrational energy levels of H_2O computed with basis sets truncated at $P_{\max} = 12$ and $P_{\max} = 20$, using two different coordinate systems: (1) a fixed linear transformation of valence coordinates, and (2) normalizing-flow coordinates optimized for H_2S with $P_{\max} = 12$. The normalizing-flow coordinates were fine-tuned using the transformation discussed in subsection 3.3.2, with six parameters, one multiplicative and one additive parameter for each coordinate. The fine-tuning was performed only for the smallest truncated basis set, $P_{\max} = 12$, and subsequently transferred without reoptimization to the larger truncation $P_{\max} = 20$ (see subsection 3.3.4). The results in Fig. 3.5 show that normalizing-flow coordinates transferred from H_2S to H_2O yield more accurate predictions for all 100 vibrational energy levels compared to the valence coordinates, for the same $P_{\max}$. This demonstrates the utility of knowledge transfer between the Hamiltonians of chemically related molecules H_2S and H_2O. The observed improvement is expected, given the morphological similarity of their PESs, which also underlies the general effectiveness of valence coordinates in both systems.

While H_2S and H_2O exhibit similar structures in their lowest-energy vibrational states, this resemblance breaks down at higher excitations of the bending mode. Notably, H_2O begins to sample quasilinear configurations within the first 100 vibrational states, unlike H_2S. These are exactly the states in Fig. 3.5 for which the convergence deteriorates when using transferred normalizing-flow coordinates. Despite this important difference in high-energy behaviour, the normalizing-flow coordinates optimized for H_2S still provide a superior representation of the bending motion in H_2O compared to valence coordiantes, within the same truncated basis set.

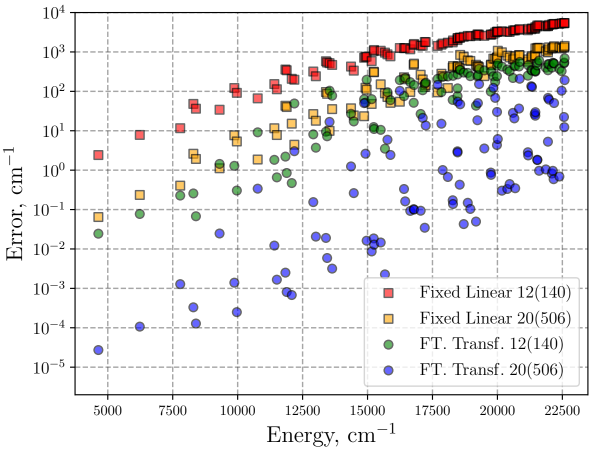

Figure 3.5: Convergence of the first 100 energy levels of H_2O for $P_{\max} = 12$ (140 basis functions) and $P_{\max} = 20$ (506) using fixed linear transformed valence coordinates (red and orange squares) compared to fine-tuned (FT.) normalizing-flow coordinates transferred from H_2S calculations, optimized for $P_{\max} = 12$, for $P_{\max} = 12$ and $P_{\max} = 20$ (green and blue circles). The vertical axis represents the calculated energy error, $E_i - E_i^{(\mathrm{Ref})}$, while the horizontal axis shows the reference energy, $E_i^{(\mathrm{Ref})}$. The converged energy values for H_2O were computed using optimized normalizing-flow coordinates with $P_{\max} = 24$.

3.4 Conclusion

The normalizing-flow approach offers a versatile framework for optimizing vibrational coordinates in variational calculations of molecular vibrational energies. In previous work, we showed that optimized normalizing-flow coordinates produce physically meaningful transformations that enhance the separability of the Hamiltonian when expressed in a direct-product basis [113].

A key question is whether the normalizing-flow method can learn intrinsic coordinates that not only improve computational performance for individual molecules, but that also generalize across chemically related systems. In this work, we took an initial step toward addressing this question.

We demonstrated that optimizing vibrational coordinates using normalizing flows can improve the accuracy of variationally computed vibrational energies by several orders of magnitude. Furthermore, these coordinates enable much faster basis-set convergence compared with both fixed linear and optimized linear transformations of valence coordinates.

Although the optimized coordinates are tailored to a specific molecule, basis-set truncation, and target set of eigenstates, our results show that this specificity does not preclude generalization. On the contrary, we find that the same coordinates often remain effective across different basis-set truncations, isotopologues, and even across molecules with similar structural motifs.

This findings suggest that normalizing flows may uncover coordinate systems that reflect fundamental features of molecular motion. While these coordinates are optimized for spectral accuracy, their capacity to generalize indicates a deeper connection to the intrinsic vibrational coordinates long sought in theoretical chemistry. Thus, variational coordinate optimization with normalizing flows may serve not only as a practical computational tool, but also as means of generating physically motivated representations of vibrational motion.

Looking ahead, several promising directions emerge. One is to develop a generalized, trainable coordinate mappings that embed molecular descriptors or parameters directly into the model, thereby enabling automatic adaptation across different chemical systems. Another is to extend the current framework to larger molecular systems. At present, the use of dense Gaussian quadrature grids imposes a practical limit on system size. To address this, we are investigating Monte Carlo integration as an alternative. While it introduces greater integration errors, it offers a viable path to higher-dimensional systems where traditional quadratures become unfeasible.

Finally, our current architecture does not yet incorporate molecular symmetry, which is known to improve efficiency in conventional vibrational treatments [119]. Enforsing symmetry in the coordinate transformation and basis function could provide similar benefits for normalizing-flow coordinates. However, imposing symmetry constraints may reduce the flexibility of the neural network, presenting a tradeoff between generalization and effectiveness.

4 Inducing dense sets via composition with strictly monotonic functions[3]

Abraham Maslow, an American psychologist of the 20th century, popularized the quote: "if all you have is a hammer, everything looks like a nail". With this quote, he encapsulated the essence of the law of the instrument, a concept also associated with the American philosopher Abraham Kaplan. The central idea of this law is that the way one approaches problems is heavily influenced by the instruments one has available.

Continuing with the metaphor, normalizing flows—invertible transformations to enhance expressivity— have been my personal hammer throughout my doctoral studies. The long thinking sessions ramificated the application of invertible transformations into other approximation problems.

In this chapter, we present a novel approach to construct flexible dense sets in the space of continuous functions. These sets are constructed using an underlying dense set and composing them with a strictly monotone transformation, i. e., a normalizing flow. This transformation acts as a non-singular change of coordinates. In particular, we explore the composite set of algebraic polynomials, which we define as the induced set. We provide a finite bound for the degree of the induced set that can produce arbitrary precision. In contrast to traditional convergence theorems where the approximation error *approaches* zero with the polynomial degree, this result provides a finite threshold for the obtention of arbitrarily precise approximations. We connect this theory to the solution of the Schrödinger equations by using this technique to approximate potential energy surfaces.

4.1 Introduction

The discovery of basis functions for functional spaces has been a cornerstone in the development of both mathematical theory and numerical analysis, with far-reaching applications across the natural and engineering sciences. Since the introduction of Fourier series, a vast array of basis functions has been developed for various function spaces, and their study remains a fundamental part of modern mathematical curricula.

Coordinate transformations have played a pivotal role in the theory of basis functions. A classic example is the Chebyshev polynomials, which are defined *via* a change of variables from Fourier cosine functions and were instrumental in addressing Runge's phenomenon in polynomial interpolation [120]. Over time, the importance of coordinate transformations in numerical approximation became increasingly apparent, particularly in resolving singularities at boundaries and corners of domains.

In the recent years, coordinate transformations have been used to construct new basis sets in L^2 [75]. These sets can be constructed by composing a known orthonormal basis set $\{\phi_n\}_{n=0}^{\infty}$ with an invertible mapping g,

$$\gamma_n(x) = \phi_n(g(x))\sqrt{\left|\det \nabla_x g(x)\right|},$$

---

[3] This chapter is based on the manuscript under preparation "Inducing dense sets *via* composition of polynomials with strictly monotonic functions" by Álvaro Fernández Corral, Yahya Saleh, and Armin Iske. My contributions to this work were the conceptualization, preparation and formalization of the theory, production of the numerical results that support the theoretical results, generation of the figures, and writing of the manuscript.

where the determinant of the Jacobian accounts for the volume change under the transformation and ensures the orthonormality of the induced basis set $\left\{\gamma_n\right\}_{n=0}^{\infty}$. This induced set was then used to approximate the solutions of partial differential equations (PDEs), particularly in quantum mechanics [95]. Applications in the fields of nuclear motion theory [92, 113, 121] and condensed matter physics [114, 122] demonstrated state-of-the-art results and orders of magnitude improvement in accuracy and convergence order.

Traditionally, coordinate transformations in numerical analysis have been manually designed by experts to suit specific problems. However, this paradigm is shifting with the emergence of *normalizing flows*, i. e., optimizable invertible coordinate transformations developed for modeling probability densities [123]. In its standard formulation, a normalizing flow is a parameterized change of variables g_θ that enhances the expressiveness of a base probability distribution in order to approximate a target distribution. Normalizing flows have become a fundamental tool in generative machine learning, widely used for modeling images, audio, and language. We borrow from this community the idea of using an invertible transformation to enhance the expressivity, and parametrize it with a numerical invertible transformation.

In this work, we study the approximation of continuous functions on compact domains in $\mathbb{R}$ using the linear span of monomials composed with strictly monotonic continuous transformations. We prove that for every set $\{\phi_i\}_{i=0}^{\infty}$ that is dense in the space of continuous functions on a compact domain, the composition set $\left\{\phi_i \circ h\right\}_{i=0}^{\infty}$ remains dense under the unique restriction that h is strictly monotonic. In particular, this condition applies to the set of monomials, creating the *induced set* $\left\{x^i \circ h\right\}_{i=0}^{\infty}$. Further, for any continuous target function exhibiting N local extrema, we demonstrate the existence of a strictly monotonic continuous function h such that the linear span of the set $\left\{x^i \circ h\right\}_{i=0}^{N+1}$ yields an approximation of the function that is accurate up to arbitrary precision.

This result can be viewed as a generalization of the classical fact that an exact representation of a function by a finite set of monomials is only possible if the function is itself a polynomial. Here, the expressivity of the set of monomials is enhanced by a coordinate transformation, enabling finite-dimensional representations for a much broader class of continuous functions.

To support our theoretical findings, we provide numerical experiments where the strictly monotonic transformation is parameterized by an invertible residual neural network (iResNet) [76]. By construnction, the iResNet produces a smooth, analytic and bi-Lipschitz transformation h_θ, implying that a vanishing first-order derivative cannot be created for any value of the parameters, which plays a crucial role in our theorems.

4.2 Notation and Background

Let $\Omega \subset \mathbb{R}$ be a compact and connected set. We denote by $C(\Omega)$ the space of continuous functions on Ω. All the functions in this manuscript are assumed to be real-valued. Unless otherwise stated, all sets considered in this study are assumed to be compact and simply connected.

By the Weierstrass approximation theorem, the set of polynomials $\Pi = \operatorname{span}\left(\{x^i\}_{i=0}^{\infty}\right)$ is dense in $C(\Omega)$ with respect to the uniform norm. That is, for any $f \in C(\Omega)$ and $\varepsilon > 0$, there exists an integer $N \in \mathbb{N}$ and a polynomial

$$p_N(x) = \sum_{i=0}^{N-1} a_i x^i$$

such that

$$\sup_{x \in \Omega} \left| f(x) - p_N(x) \right| < \varepsilon.$$

The integer N is called the degree of the approximating polynomial p_N.

An exact representation of a function $f \in C(\Omega)$ by a polynomial of finite degree is only possible if f is itself a polynomial. In general, polynomial approximations converge to f asymptotically as $N \to \infty$, with the convergence rate governed by the smoothness of f. Classical results such as Jackson's inequality provide quantitative bounds on this approximation error in terms of the function's regularity.

The Müntz-Szász theorem extends the number of available dense sets in $C(\Omega)$, under the assumption $\Omega = [a, b]$ and $b > a > 0$. Then, the power set defined by $\Psi_\Lambda = \operatorname{span}\left(\{x^{\lambda_i}\}_{i=0}^{\infty}\right)$ is dense in $C(\Omega)$ if $\lambda_0 = 0$ and $\sum_{i=1}^{\infty} 1/\lambda_i = \infty$, with $\Lambda = \{\lambda_i\}_{i=0}^{\infty} \subseteq \mathbb{N}$. Extensions of the Müntz-Szász theorem allow for the creation on dense sets of C in any compact interval of the real line. Note that in particular the polynomial set satisfies the Müntz-Szász theorem, i. e., $\Pi = \Psi_\Lambda$ if $\Lambda = \mathbb{N}$.

4.3 Approximation *via* polynomials enhanced by invertible transformations

In this section, we investigate the improved approximation capabilities achieved by composing dense sets with continuous, strictly monotonic functions. We begin by establishing the theoretical guarantees for the density of such composed sets in the space of continuous functions.

Theorem 4.1. *Let* $\Phi := \operatorname{span}\left(\{\phi_i\}_{i=0}^{\infty}\right)$ *be dense in* $C(\Omega_h)$ *with respect to the uniform norm. Let* $h : \Omega \to \mathbb{R}$ *be a continuous, and strictly monotonic function and denote its image by* $\Omega_h := h(\Omega)$. *Then the set*

$$\Phi_h := \operatorname{span}\left(\{\phi_i \circ h\}_{i=0}^{\infty}\right)$$

is dense in $C(\Omega)$ *with respect to the uniform norm.*

Proof. Let $f \in C(\Omega)$ be arbitrary. Since $h : \Omega \to \mathbb{R}$ is continuous and strictly monotonic, it is injective and thus admits a well-defined inverse $h^{-1} : \Omega_h \to \Omega$. Moreover, by standard results [124], h^{-1} is unique, continuous and strictly monotonic.

The composition of continuous functions is continuous [124, 125], so $f \circ h^{-1} \in C(\Omega_h)$. By the density of $\Phi \subset C(\Omega_h)$, for any $\varepsilon > 0$, there exists a function $g \in \Phi$ such that

$$\sup_{x \in \Omega_h} \left| f \circ h^{-1}(x) - g(x) \right| < \varepsilon.$$

Define the function $\tilde{g} := g \circ h$. Since $g \in \Phi = \operatorname{span}\left(\{\phi_i\}_{i=0}^{\infty}\right)$, it follows that $\tilde{g} \in \Phi_h = \operatorname{span}\left(\{\phi_i \circ h\}_{i=0}^{\infty}\right)$. For all $y \in \Omega$ and $x = h(y)$, we then have

$$\left| f(y) - \tilde{g}(y) \right| = \left| f(y) - g(h(y)) \right| = \left| f \circ h^{-1}(x) - g(x) \right| < \varepsilon,$$

since $\Omega_h = h(\Omega)$. Therefore,

$$\sup_{y\in\Omega} |f(y) - \tilde{g}(y)| < \varepsilon.$$

This shows that any $f \in C(\Omega)$ can be approximated arbitrarily closely with respect to the uniform norm by functions in Φ_h, and thus Φ_h is dense in $C(\Omega)$. □

Theorem 4.1 assumes the image of the strictly monotonic mapping h to be the domain of the dense set Φ. This imposes a strong constrain in the possibilities for the transformation h. This limitation can be overcome by the use of sets that are dense on any compact subset of the real line. In particular, the set of polynomials $\Pi = \text{span}\left(\left\{x^i\right\}_{i=0}^{\infty}\right)$ satisfies this condition. This motivates the definition of the induced set $\Pi_h := \text{span}\left(\left\{x^i \circ h\right\}_{i=0}^{\infty}\right) = \text{span}\left(\left\{h^i\right\}_{i=0}^{\infty}\right)$, which is dense for every domain and image of h.

We note that, in the special case in which Φ satisfies the Müntz-Szász theorem, Theorem 4.1 resembles a density result reported by Jaming and Simon [126]. Specifically, the authors show that the set $\Phi_h = \text{span}\left(\{h^{\lambda_i}\}_{i=0}^{\infty}\right)$— where the set of powers $\Lambda = \{\lambda_i\}_{i=0}^{\infty}$ satisfy the Müntz-Szász conditions— is dense in the space of continuous functions only if h is monotonic and its first and second-order derivative do not vanish simultaneously [126, Proposition 3.1]. In contrast, our result does not require h to be differentiable.

In what follows, we restrict our analysis to the study of the approximating properties of the induced set Π_h. For any continuous function with finitely-many local extrema, we show that the degree of the induced polynomial that approximates f to arbitrary precision is finite. To prove this, we first introduce the following lemma.

Lemma 4.1. *Let f be a monotonic, continuous functions $f : \Omega_f \to \mathbb{R}$ and let g be a strictly monotonic, continuous functions $g : \Omega_g \to \mathbb{R}$. Let $f(\Omega_f) = g(\Omega_g) =: I$. Then, for every $\varepsilon > 0$, there exists a strictly monotonic, continuous, surjective function $h : \Omega_f \to \Omega_g$, such that*

$$\sup_{y\in\Omega_f} \left|f - g \circ h\right| < \varepsilon.$$

Proof. Using the density of strictly monotonic functions in the space of monotonic functions, there exists a continuous strictly monotonic function $\hat{f}$ such that, for every $\varepsilon > 0$,

$$\sup_{y\in\Omega_f} \left|f - \hat{f}\right| < \varepsilon.$$

Since g is strictly monotonic and continuous functions, its inverse function g^{-1} exists and is unique, continuous and strictly monotonic [124]. The function $h = g^{-1} \circ \hat{f}$ satisfies the conditions of the lemma. □

Having shown that two functions with the same image on a compact domain can be mapped onto each other *via* a non-singular coordinate transformation, we are ready to prove the main result of the manuscript.

Theorem 4.2. *Let $f : \Omega_f \to \mathbb{R}$ be a contiuous function defined on a compact interval $\Omega_f \subset \mathbb{R}$. Suppose that f has exactly $M \in \mathbb{N}$ local extrema on Ω_f. Then, there exists*

- *a compact interval $\Omega_h \subset \mathbb{R}$,*

- *a strictly monotonic, continuous, and surjective map* $h : \Omega_f \longrightarrow \Omega_h$*,*
- *and coefficients* $a_0, a_1, \dots, a_{M+1} \in \mathbb{R}$*,*

such that, for any $\varepsilon > 0$*,*

$$\sup_{x\in\Omega_f} \left| f(x) - \sum_{i=0}^{M+1} a_i h^i(x) \right| < \varepsilon. \tag{4.1}$$

Moreover, if f *has no saddle points and its extrema are sets of measure zero, then*

$$f(x) = \sum_{i=0}^{M+1} a_i h^i(x) \tag{4.2}$$

for all $x \in \Omega_f$*.*

Proof. Let $\Omega_f = [x_0, x_{M+1}]$ and denote by $x_1, \dots, x_M$ the M local extrema of f such that $x_0 < x_1 < x_2 < \cdots < x_{M+1}$. If a local extrema has measure non-zero, choose any point of the set arbitrarily. Set $I_0 := [x_0, x_1]$ and $I_i := (x_i, x_{i+1}]$ for all $1 \leq i \leq M$. Note that the set

$$P_f := \{x_0, \dots, x_{M+1}\}$$

partitions Ω_f such that the restrictions $f|_{I_i}$ are, in general, monotonic.

Denote the functional extreme values by $f_i = f(x_i)$ for $i = 0, \dots, M+1$. As these values are evaluations at extreme points, they follow the recurrent relation $\pm(-1)^i(f_{i+1} - f_i) > 0$ for all $i = 0, \dots, M$ for either positive or negative sign. Effectively, this means that if $f_i > f_{i-1}$, then $f_i > f_{i+1}$ as well.

Let $\Omega_h = [y_0, y_{M+1}]$. Let $p(y) = \sum_{i=0}^{M+1} a_i y^i$ be a polynomial of degree $M+1$, $y \in \Omega_h$, $a_i \in \mathbb{R}$ for all i. The polynomial coefficients $\{a_i\}_{i=0}^{M+1}$ and a set of points $\{y_i\}_{i=0}^{M+1}$, $y_0 < y_1 < y_2 < \cdots < y_{M+1}$, are chosen to satisfy

$$\begin{cases} p\left(y_i\right) = f_i, & i = 0, \dots, M+1 \\ y_i \text{ is an extremum of } p \text{ for} & i = 1, \dots, M. \end{cases} \tag{4.3}$$

A proof of the existence of a solution to this problem is provided in subsection 4.7.1.

Set $J_0 := [y_0, y_1]$ and $J_i := (y_i, y_{i+1}]$ for all $1 \leq i \leq M$. As p is of degree $M+1$ and satisfies (4.3), the set $P_p := \{y_0, \dots, y_{M+1}\}$ partitions Ω_h such that the restrictions $p|_{J_i}$ are strictly monotonic. In addition, this construction imposes that the images of the restrictions are shared $f(I_i) = p(J_i)$ for all $i = 0, \dots, M$. By Lemma 4.1, for every positive constant $\alpha_i > 0$, there exist a continuous, strictly monotonic, surjective transformation $h_i : I_i \longrightarrow J_i$ such that

$$\sup_{x\in I_i} \left| f(x) - p \circ h_i(x) \right| < \alpha_i, \quad i = 0, \dots, M.$$

Define the global mapping $h : \Omega_f \longrightarrow \Omega_h$ piecewise by

$$h(x) := h_i(x), \quad \text{for } x \in I_i, \quad i = 0, \dots, M.$$

Since each h_i is continuous, strictly monotonic and surjective, and the intervals I_i only meet at endpoints, the concatenation defines a continuous, piecewise strictly monotonic, surjective map $h : \Omega_f \to \Omega_h$.

It follows that, for every positive constant $\varepsilon > 0$, there exists a strictly monotonic, surjective function $h : \Omega_f \to \Omega_h$ that satisfies

$$\sup_{x\in\Omega_f} \left| f(x) - p \circ h(x) \right| = \max_i \left(\sup_{x\in I_i} \left| f(x) - p \circ h_i(x) \right| \right) < \varepsilon,$$

completing the proof of (4.1).

Last, if f has no saddle points and the set of local extrema has measure zero, every restriction $f|_{I_i}$ is strictly monotonic. It follows from the proof of Lemma 4.1 that in this case $f(x) = p \circ h_i(x)$ is achievable for all $x \in I_i$ and every $i = 0, \ldots, M$. Therefore, the approximation achieves exactness in this case, proving (4.2). □

Theorem 4.2 establishes the existence of an invertible transformation h and a corresponding induced set that allows a finite dimensional arbitrary-precision approximation of a broad class of functions. This finite-degree threshold constrasts convergence theorems, such as Jackson's inequality, where the approximation error diminishes as the degree of the polynomial increases. Here, the error depends only on of h.

To gain intuition on the approximation within the induced set Π_h, we consider the case of functions with a single local extremum. In such a special case, the mapping h can be constructed analytically.

Example 4.1. Let $f \in C(\Omega)$ be a continuous function with a single local extremum at x_0. By Theorem 4.2, there exists a strictly monotonic function h and real coefficients a_0, a_1, a_2 such that

$$\left\| f - \sum_{i=0}^{2} a_i h(x)^i \right\| < \epsilon, \qquad \forall \epsilon > 0.$$

To construct such an approximation, set $a_0 = f(x_0)$ and define the remainder $\hat{f} = f - a_0$. Since $\hat{f}(x_0) = 0$, its image is either nonnegative or nonpositive on Ω. The first-order term is unnecessary to capture the shape of a function with a single extremum, and therefore $a_1 = 0$. Thus the approximation problem reduces to representing $\hat{f}$ using $a_2 h(x)^2$. The second-order coefficient is chosen so that $\hat{f}/a_2 \geq 0$ as $a_2 \in \{+1, -1\}$.

This leads to the representation

$$\hat{h}(x) = \operatorname{sign}(x - x_0) \sqrt{|\hat{f}(x)|}, \tag{4.4}$$

which is continuous since the sign discontinuity occurs exactly at x_0, where $\hat{f}(x_0) = 0$.

In general, $\hat{h}$ needs not be strictly monotonic (as $\hat{f}$ needs not neither). Because strictly monotonic functions are dense in the space of monotonic functions, for any $\varepsilon > 0$ there exists a strictly monotonic function h satisfying

$$\sup_{x\in\Omega} |h(x) - \hat{h}(x)| < \varepsilon.$$

Therefore, the composition of h with the second-degree polynomial with coefficients $(a_0, 0, a_2)$, provides an approximation of f to arbitrary precision.

4.4 Numerical results

In this section, we present numerical experiments that support our theoretical results. For a given target function f, we construct approximations of the form

$$\hat{f}_\Theta = \sum_{i=0}^{N-1} a_i h_\Theta^i(x),$$

where the nonlinear change of coordinates h_Θ was modelled *via* an iResNet. By construction iResNets are smooth bi-Lipschitz transformations and hence satisfy the assumptions on h in Theorem 4.1 and Theorem 4.2. We constructed the iResNet using 15 residual blocks, each composed of 2 hidden layers of 8 units each and that use LipSwish as activation function.

We compare the optimization results with standard polynomial approximations,

$$\hat{f} = \sum_{i=0}^{N-1} a_i x^i.$$

We demonstrate that optimizing the nonlinear change of coordinates h enhances approximation of a broad class of continuous target functions with different number of local extrema, as described in Lemma 4.1 and Theorem 4.2.

4.4.1 Approximating continuous functions with a single extrema

In Example 4.1, we established that any continuous function with a single local extrema can be approximated by a polynomial of degree 2 composed with a strictly monotonic continuous transformation h. We now provide numerical examples illustrating this result.

Consider the target function

$$f(x) = \exp(x) + \exp(-x), \qquad \Omega = [-10, 10],$$

which has a single extremum at $x_0 = 0$ and no saddle points.

Since the minimum occurs at $x_0 = 0$, we set $a_0 = f(x_0) = 2$, and since the linear term is unnecessary for capturing the single-extremum structure, we take $a_1 = 0$. Because the remainder $\hat{f} = f - a_0$ is nonnegative, we set $a_2 = 1$. The corresponding transformation h is given analytically by (4.4), that is

$$h(x) = \operatorname{sign}(x)\sqrt{\exp(x) + \exp(-x) - 2}.$$

In addition to this closed-form construction, we also compute a numerical approximation of the form

$$\hat{f}(x) = \sum_{i=0}^{2} a_i \, h_{\Theta^*}(x)^i,$$

where h_Θ is parametrized by an iResNet. Given a training set of equidistant points $D = \{x_p\}_{p=0}^{399}$, the neural network parameters Θ^* are obtained from the optimization problem

$$\Theta^* = \arg\min_{\Theta} \left(\sup_{x \in D} \left| f(x) - \sum_{i=0}^{2} a_i \, h_\Theta(x)^i \right| \right). \tag{4.5}$$

This nonlinear optimization was solved using using the `Optax` [115] implementation of Adam, a gradient-based stochastic optimization algorithm.

As a benchmark, we also optimized a second-degree polynomial in the original input space by solving

$$\mathbf{Xa} = \mathbf{f},$$

with $\mathbf{a} = [a_0, a_1, a_2]$, $\mathbf{f} = [f(x_0), \ldots, f(x_{399})]$, and $\mathbf{X}^{ip} = x_p^i$.

The results are reported in Fig. 4.1. Interestingly, the learned transformation h_Θ closely resembles the analytically derived mapping h, despite no explicit constraint enforcing this behavior. For validation, we computed the RMSE over a grid of 5001 equidistant points, obtaining a value of 13.506. Considering that f attains values on the order of 2×10^4, this error is remarkably small. In comparison, achieving a similar accuracy with direct polynomial fitting requires a polynomial of degree at least 10.

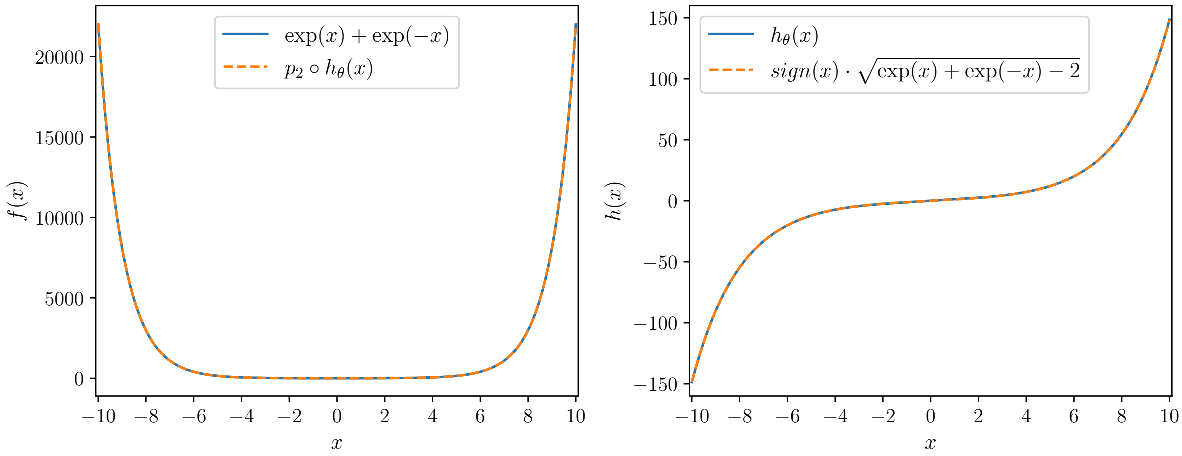

Figure 4.1: *Left:* The target function $f(x) = \exp(x) + \exp(-x)$ and the fitted second-order power series expansion $p_2(x) = x^2 + 2$ composed with an invertible bi-Lipschitz function h_θ, parametrized by an iResNet with parameters θ. The RMSE over a grid of 1000 equidistant points in the domain $[-10., 10.]$ obtained for the fit is 10.917. *Right:* The invertible function that resulted as the solution of the optimization and its comparison with the expected function $\hat{h}(x) = \text{sign}(x) \cdot \sqrt{f(x) - 2}$.

4.4.2 Approximating non-differentiable functions

Consider the non-differentiable target function

$$c(x) = \begin{cases} 1 - (x-1)^2, & x > 0 \\ \arctan(-x), & x \leq 0 \end{cases}$$

for $x \in [-3, 3]$. This function is continuous over its domain, as both defining functions are continuous and their values at the discontinuity point match. The function contains two strict local minima in its domain, placed at the discontinuity point $x = 0$ and at $x = 1$. Therefore, an arbitrarily-accurate approximation using third-degree polynomial composed with the non-singular coordinate transformation h is possible.

As no analytical choice of $\{a_i\}_{i=0}^3$ is available for this case, the objective of the optimization is to be find both the optimal a_i^* and Θ^*. Given a training set of equidistant points

$D = \{x_p\}_{p=0}^{300}$, the linear coefficients a_i^* and network parameters Θ^* are obtained from the optimization problem

$$a_i^*, \Theta^* = \arg\min_{a_i, \Theta} \left(\sup_{x \in D} \left| f(x) - \sum_{i=0}^{2} a_i \, h_\Theta(x)^i \right| \right). \tag{4.6}$$

This nonlinear problem was solved in an alternating optimization scheme: the nonlinear parameters Θ were updated with Adam, while the linear coefficients a_i were determined by solving the least-squares system

$$\sum_{i=0}^{2} a_i h_\Theta(x)^i \approx f(x), \qquad x \in D,$$

which amounts to computing the pseudoinverse of the matrix

$$X_\Theta^{ip} = h_\Theta(x_p)^i, \qquad i = 0, 1, 2, \;\; p = 0, \ldots, 300.$$

This pseudoinversion guarantees that the coefficients a_i are optimal for every parameter set Θ.

We illustrate the result of this numerical experiment in Fig. 4.2. The approximation error is larger in the neighborhood of $x = 0$, as $p_3 \circ h$ is smooth and this is the discontinuity point of f.

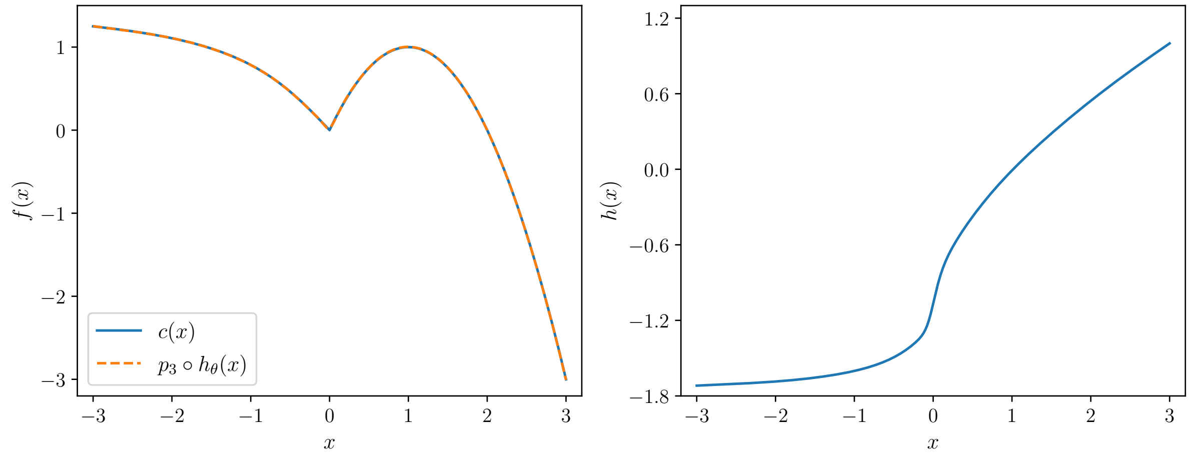

Figure 4.2: *Left:* The target function and the fitted second-order expansion composed with an invertible function. The RMSE over a grid of 5001 equidistant points in the domain $[-3, 3]$ for the fit is 0.19. *Right:* The invertible function that resulted as the solution of the optimization. The discontinuity in the derivatives at $x = 0$ produces an invertible transformation h whose derivative tends to infinity at that point.

For validation, a polynomial was fitted on the training points and both approaches were evaluated over an equidistant grid of 5001 points in the domain. In the validation set, the induced set achieved a RMSE of 0.19 and a maximum absolute error 0.028. In contrast, the best direct polynomial fit, using a polynomial of degree 80, yielded a RMSE of 0.490 and a maximum absolute deviation of 0.063. Both error metrics for the direct fit are approximately an order of magnitude larger than those of the induced set, highlighting its superior performance in approximating non-smooth functions.

4.4.3 Approximating functions with non-strict local extrema

Last, we demonstrate the effectiveness of the induced set approximation for continuous function that contain local extrema that are not strict. In this example, the target function is given by

$$f(x) = \begin{cases} \exp\left(-\frac{1}{(x-1)^2}\right), & |x| > 1 \\ 0, & |x| \leq 1 \end{cases}, \tag{4.7}$$

and is defined in the domain $\Omega = [-4, 4]$. f is continous and contains an interval of constant values $J = [-1, 1]$ which constitutes a local minimum. No other extrema is present in the domain of f. For this reason, a second-degree polynomial composed with an invertible function h can be used to approximate it.

Once more, we use the strategy of Example 4.1. We select the second-degree polynomial p_2 with coefficients $a_0 = f|_J = 0$, $a_1 = 0$ and $a_2 = 1$. This polynomial is then composed with an invertible function h. The parameters of the neural network were trained using a grid of 1000 equidistant points over the function domain in the same fashion as in (4.5). The results of this approximation are shown in Fig. 4.3.

The results were validated by comparison to a standard polynomial fit in the training points and evaluated over a validation grid of 5000 points equidistant points. The obtained RMSE for the induced set fit in the validation set was $9.40 \cdot 10^{-4}$, and the maximum absolute error of this fit was $2.67 \cdot 10^{-3}$. The error of the approximation increases as one moves away from its center. This is attributed to the contraction of the flat region into a set of near-zero measure by the invertible transformation, as reflected in the flatness of the learned function h around this point.

To replicate the accuracy of the composed set using a polynomial fit, a polynomial of degree 40 was required. This approximation finds a RMSE = $9.16 \cdot 10^{-4}$ in the validation grid and a maximum deviation of $2.65 \cdot 10^{-3}$.

4.5 Applications in multidimensional problems

4.5.1 The multidimensional induced set

While the one-dimensional theory for the finite degree of the induced set approximation does not extend straightforwardly to multidimensional problems, the utility of this construction can be nevertheless layed out.

Let the target function $f \in C(\Omega)$ be defined in a multidimensional domain $\Omega \subset \mathbb{R}^D$. To approximate such a function, a multidimensional polynomial can be constructed using direct products of the one-dimensional monomials. Let the coordinates of the system be denoted by x_α, $\alpha = 1, \ldots, D$. Then, the multidimensional polynomial is constructed as a linear combination of the functions in the direct product set $\left\{x_1^{i_1} x_2^{i_2} \ldots x_D^{i_D}\right\}_{i_0, i_1, \ldots i_D = 0}^{\infty}$, which is dense in $C(\Omega)$. In practice, only polynomials up to some finite degree $N \in \mathbb{N}$ are used. The total number of functions in the set— and thus, of linear coefficients to fit— grows roughly as $(N+1)^D$. This exponential scaling with the dimensionality is known as the curse of dimensionality.

Identically to the one-dimensional case, the approximation capactiy of multidimensional polynomials can be enhanced by the composition with an invertible function. Let $h : \Omega \to \Omega_h \subset \mathbb{R}^D$ be a continuous, strictly monotone, surjective mapping. Define $q_k =$

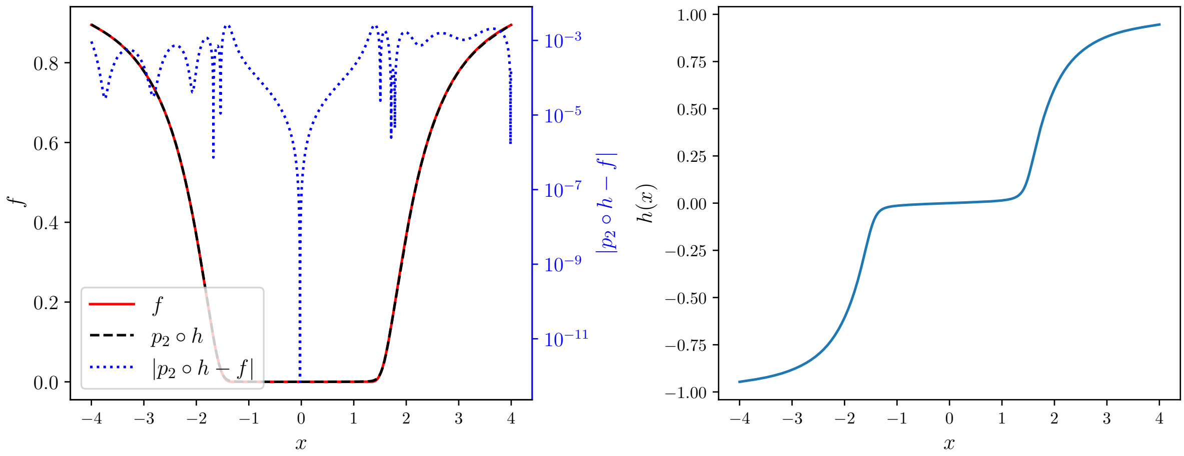

Figure 4.3: *Left:* The target function (solid red) and the fitted second-order expansion composed with an invertible function (dashed black). The error of the fit is also ploted (dotted blue) for a visualization of the points where the fitting achieves higher and lower accuracy. The RMSE of the fit on the validation grid of 5000 equidistant points in the domain $[-4, 4]$ is $9.40 \cdot 10^{-4}$ and the maximum error is of $2.67 \cdot 10^{-3}$. *Right:* The obtained invertible transformation for the fit. In the interval of constant values, h remains close to flat, contracting the interval to a set of measure near-zero.

$h_k(\mathbf{x})$, with $k = 1, \ldots D$, such that h performs a non-singular coordinate transformation from $\mathbf{x}$ to $\mathbf{q}$. The resulting induced set takes the direct product form

$$\Pi_h^D = \operatorname{span}\left(\left\{q_1^{i_1} q_2^{i_2} \ldots q_D^{i_D}\right\}_{i_0, i_1, \ldots, i_D = 0}^{\infty}\right). \tag{4.8}$$

We use Theorem 4.1 to recall that the multidimensional induced set remains dense in $C(\Omega)$. This is possible since at no point of the theorem or its proof we used the fact that the domains where one-dimensional.

The one-dimensional examples demonstrate that the induced set significantly reduces the required number of functions necessary to achieve high-accuracy. If this trend extends to the multidimensional case, it would mitigate the curse of dimensionality by reducing the number of functions per dimension.

We illustrate the application of the multidimensional induced set in the approximation of continuous functions over multidimensional domains and compare its performance against standard polynomial fits. Let the target function be $f(x, y) = \arctan(x) \cdot \arctan(y)$ be defined in the two-dimensional domain $\Omega = [-4, 4] \times [-4, 4]$. f is continuous in Ω. To approximate f, we used a polynomial of second-degree denoted by p_2 composed with a strictly monotone function h, i. e., $p_2 \circ h$. This construction results in a set of 6 functions. The parameters of the numerical transformation were trained using a direct product grid of 20 equidistant points per dimension, and the linear coefficients of the expansion were found by pseudo-inversion of the evaluation matrix in the training points. The results of this numerical simulation are shown in Fig. 4.4.

For comparison, we fitted a standard polynomial on the training points and evaluated the accuracy of both methods in a validation grid of 100 equidistant points per dimension. Over the validation points, the induced set created an approximation with a RMSE of 0.031 and a maximum error of 0.0011. The best polynomial fit achievable was that of degree 13, denoted by p_{13} and containing 105 functions. By simple inspection of the figure, this polynomial fit performed significantly worse than the induced set. The maximum error of the polynomial fit was 0.163, that is two orders of magnitude worse than the induced set, and the RMSE was 2.339. These results show that the induced set not only requires an order of magnitude fewer functions, but also outperforms polynomial fit by two orders of magnitude in both metrics.

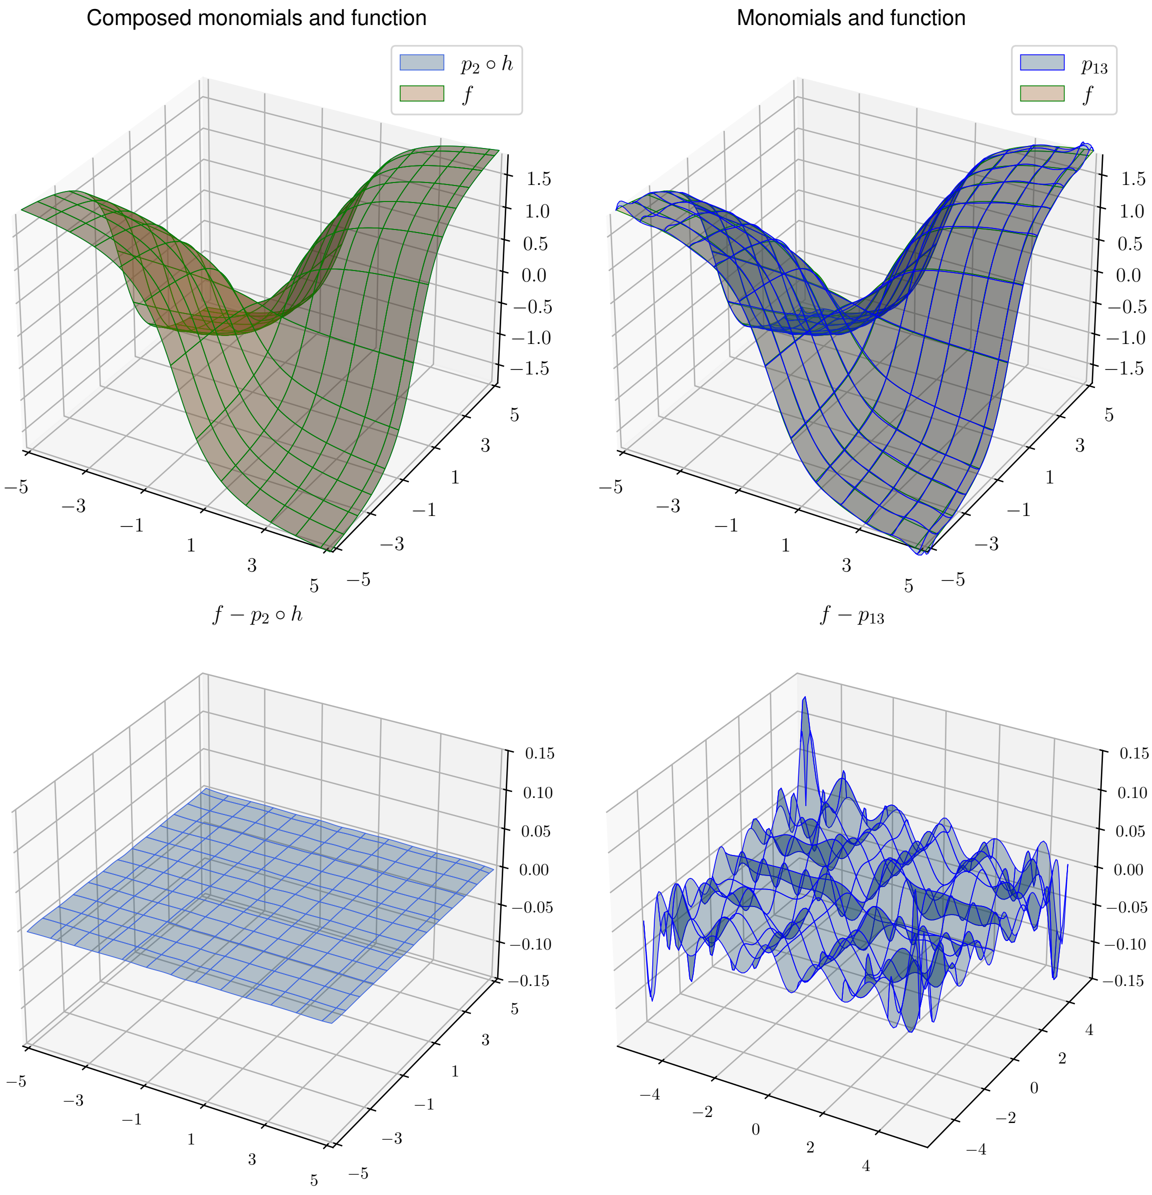

Figure 4.4: *Top left:* The target function and the fitted second-order expansion composed with an invertible function. The RMSE of the fit on the validation points was 0.031. *Bottom left:* The difference between the target function and the composed function on the validation set. The maximum difference was 0.0011. *Top right:* The target function and the fitted 13-degree polynomial. The RMSE of this fit on the validation set was 2.339. *Bottom right:* The difference between the target function and the 13-degree polynomial on the validation set. The maximum difference was 0.163.

4.5.2 Potential energy surface fitting

The ability to approximate multidimensional functions with high accuracy using low-degree polynomials has broad implications in fields where functional approximation are required but hindered by the curse of dimensionality. One such field is the fitting of potential energy surface (PES) in molecular physics. We introduce the meaning and importance of this function shortly.

When solving the Schrödinger equation for molecular systems, the Born-Oppenheimer approximation is commonly employed to decouple the electronic and nuclear degrees of freedom. This separation is justified by the great difference in timescales of electronic and nuclear motions. Under this approximation, the electronic Schrödinger equation is first solved at fixed nuclear geometries. The PES is the function that maps the nuclear configuration to corresponding electronic energies. Accurate representations of the PESs are key for the calculations of the spectroscopic and dynamical properties of molecules. However, generating the data to construct the PES is both difficult and computationally demanding, as each point requires of the solution of the electronic Schrödinger equation. Therefore only small sample sizes are available, and approximation techniques that present good interpolation and extrapolation properties are highly desirable.

Let f be the target PES, which is a real-valued, continuous function defined on a domain $\Omega \subseteq \mathbb{R}^D$. A standard approach to fit PESs involves expanding the function as a power series expansion in terms of functions of the coordinates of the problem $\mathbf{x}$ [127–130]. This initial transformation is typically defined as

$$y_j = g_j(x_j), \quad j = 1, \dots, D, \tag{4.9}$$

where $g_j : \Omega_j \to I_j$ is a different function per variable, $\Omega_j \subseteq \mathbb{R}$ is the domain of x_j and $I_j \subseteq \mathbb{R}$ is the image of g_j. The complete domain Ω is the Cartesian product of all Ω_j, and $\mathbf{y}$ denote the transformed variables. The choice of the function g_j is generally guided by the physical meaning of x_j and often involves few parameters.

The PES is then approximated by the power series in the variables $\mathbf{y}$, truncated at degree N

$$V(\mathbf{x}) = \sum_{i_0+\cdots+i_D=0}^{N} c_{i_0,\dots,i_D}\, y_0^{i_0}(x_0)\cdots y_D^{i_D}(x_D), \tag{4.10}$$

where $c_{i_0,\dots,i_D}$ are the coefficients of the expansion.

In the PES literature, the functions g_j are usually chosen to be strictly monotonic. However, to the best of our knowledge, no formal justification for this constraint has been provided, beyond empirical sucess. We propose that the underlying rationale can be understood through the lens of the density result presented in Theorem 4.1. Specifically, we refer to the fact that composing the multidimensional polynomials with strictly monotone transformations yields a set that remains dense in C. This deeper theoretical insight provides a principled explanation for the enhanced convergence of the power series in the transformed variables $\mathbf{y}$, particularly when the transformation encapsulates the behavior of the PES.

Alternative methods for PES fitting include the use of neural network to directly model the surface [131–133]. We refer the reader to the comprehensive review Manzhos and Carrington[105] for an in-depth discussion of these approaches.

The class of coordinate-wise strictly monotonic transformations g_j in (4.9) can be viewed as special classes of the family of more general multidimensional strictly monotone transformations $\mathbf{q} = h(\mathbf{x})$, where $h : \Omega \to \Omega_h$, with $\Omega_h \subseteq \mathbb{R}^D$. In this broader setting, the transformation is constructed jointly across all coordinates. Such strictly monotone transformation can be modelled by iResNets, which are analytic and have a non-zero determinant of the Jacobian by construction. The PES can then be approximated as a power series of the variable $\mathbf{q}$, following the same structure as (4.10). We recall that the image of h, $\Omega_h = h(\Omega)$, is arbitrary due to the density of the polynomial set in any compact domain.

To assess the approximation efficiency of the induced set relative to established approaches and the direct application of polynomials, we performed fits of the PES (in units of cm^{-1}) of the H_2S molecule. The variables of the problem, x_0, x_1, x_2 represent the two hydrogen-sulfur distances (H-S) in Å, and the interior bond angle in radians, respectively. The domain of the radial coordinates is $[0, \infty)$ and the domain of the angular coordinate is $[0, \pi]$. In practice, the domains of the radial coordinates are truncated to a finite interval $[a, b]$ $(b > a > 0)$ to focus on physically relevant configurations, producing a compact domain that now allows for the approximation *via* polynomials.

The analytic representation of the PES was computed by Azzam et al. [83] using an 8-th degree polynomial in the transformed set of coordinates

$$\begin{aligned} y_0 &= 1 - \exp(-\alpha_0(x_0 - \beta_0)) \\ y_1 &= 1 - \exp(-\alpha_1(x_1 - \beta_1)) \\ y_2 &= \cos(x_2) - \cos(\beta_2), \end{aligned} \tag{4.11}$$

where the parameters $\{\alpha_i\}_{i=0,1}$ control the width of the minima of the PES and the parameters $\{\beta_j\}_{j=0,1,2}$ specify the position of the molecular equilibrium configuration, i. e., the position of the minimum of the PES. These three transformations are strictly monotonic in the domain of the problem. While the one-sided derivatives of the $\cos$ function at the boundaries of the domain are zero, we remark that this does not contradict the definition of strict monotonicity, as it is based on the differences between functional values at any two points of the domain.

This analytical representation constitutes the ground truth for our fitting problem. The direct comparison to the original fitting errors is not possible, as the raw data for the fit was not available. However, it is notable that the prior fit required a polynomial of degree 8, together with a good physical understanding of the system to select the functions g_j. This polynomial contains 94 coefficients to be computed in total.

The parameters of the iResNet and linear coefficients were optimized in a grid consisting of a direct product of 40 points per dimension in the pysically relevant region, $\Omega = [0.9, 2.5] \times [0.9, 2.5] \times [0, \pi]$. The PES computed on this grid constitute the training data. To mitigate numerical instabilities due to the bad extrapolation of the fitted PES in regions with large values— therefore less important for the molecular dynamics—, the points with a value $V > 40000\ \mathrm{cm}^{-1}$ were taken out of the optimization. This cutoff exceeds the molecular dissociation energy, i. e., the energy that is required to break one of the bonds of the molecule, so that the long-range behaviour is captured in the set. In addition, the PES minimum was manually incorporated to the training set to guarantee a good calculation of the zero-degree term of the polynomial.

For validation, a standard polynomial in the original coordinates $\mathbf{x}$ was fitted using the training points, and the accuracy of both approaches was calculated on a denser grid consisting of the product of 100 equidistant points per dimension, created following the same logic as for the training set.

The induced set of degree 4 (with a total of 35 functions) produced a RMSE error of 5837.571 in the validation grid. This accounts for a mean error of 7 cm^{-1} over the points. The best polynomial fit was achieved for degree 18 (higher degrees led to degraded fits). This polynomial was created using a total of 1330 functions. When evaluating the polynomial fit in the validation set, a RMSE of 29915.811 was obtained, an order of magnitude larger than that of the induced set.

In summary, the induced set approach yields fitting accuracy roughly and order of magnitude better than the direct polynomial fit, while requiring almost two orders of magnitude fewer functions. This underscores the efficiency and improved convergence afforded by the induced set framework in multidimensional problems.

4.6 Conclusions

In this work, we developed a one-dimensional theory for enhancing the approximation of continuous functions using the composition of polynomials with strictly monotonic transformations, yielding the induced set. We demonstrated that this set remains dense in the space of continuous functions, while offering the flexibility to significantly decrease the number of terms in the polynomial expansion. Furthermore, we introduced a theorem that allows us to predict the minimum required degree of the induced set that is capable of achieving arbitrary accuracy. This bound is determined by the number of maxima and minima that the function has, or more generally, the number of monotonic segments of the target function. This provides a general approach for accurately estimating continous function with orders of magnitude improved efficiency compared to the direct use of polynomials. All theoretical claims were validated through numerical examples that use iResNets to model the transformations.

Building on the one-dimensional theory, we extend the framework to multidimensional problems. By composing multidimensional polynomials with strictly monotone functions, an induced set that remains dense in the space of continuous functions in multidimensional spaces can be created. The transformation h enables the polynomial to capture complex structures in high-dimensional domains, dramatically reducing the required number of basis functions. This alleviates the bottleneck associated to the exponential scaling of the curse of dimensionality. Importantly, the induced set reduces the need for problem-specific expertise, as the coordinate transformation is learned through optimization. This creates a framework that is directly applicable to a wide variety of approximation problems.

We showcased the effectiveness of this methodology by fitting the potential energy surface of the H_2S molecule. The results demonstrate that the induced set achieves a significantly higher accuracy with at least an order of magnitude fewer terms than the direct fit of a polynomial.

4.7 Appendix

4.7.1 Existence of solution of the interpolation problem with varying nodes

Theorem 4.3. *Let* $\{f_i\}_{i=0}^{M+1}$ *be a set of values satisfying* $\pm(-1)^i(f_{i+1} - f_i) > 0$ *for either positive or negative and* $f_i \in \mathbb{R}$ *for all* i*. Then, there exists a polynomial* p *of degree* $M + 1$*, and a set of* *different and sorted points* $\{y_i\}_{i=0}^{M+1}$ *(*$y_0 < y_1 < \cdots < y_{M+1}$*), such that*

$$\begin{cases} p\ (y_i) = f_i, & i = 0, \dots, M+1 \\ y_i \textit{ is an extremum of } p \textit{ for} & i = 1, \dots, M. \end{cases} \tag{4.12}$$

Proof. Let $\{y_i\}_{i=0}^{M+1}$ be a set of arbitrary different and sorted points, and let $Y = \{y_i\}_{i=1}^{M}$.

Define the derivative of the polynomial p by

$$p'(y) = c \prod_{i=1}^{M} (y - y_i), \tag{4.13}$$

where $c \in \mathbb{R}$ is an arbitrary multiplication constant.

Integrating equation (4.13), the polynomial is given by

$$p(y) = c_0 + c \int dy \ \prod_{i=1}^{M} (y - y_i), \tag{4.14}$$

where $c_0 \in \mathbb{R}$ is the constant of integration. This polynomial has degree $M + 1$. By construction of this ansatz, the nodes are extrema of the polynomial and the second condition of the theorem is satisfied.

Denote the space of polynomials of degree $M + 1$ with M local extrema by ρ_{M+1}. In other words, ρ_{M+1} is the space of polynomials of degree $M + 1$ whose derivative has non-degenerate real roots. Following their definition, every polynomial $P \in \rho_{M+1}$ can be expressed in the form of equation (4.14), by choosing the parameters c_0, c and Y.

The interpolation problem in equation (4.12) can be solved using a polynomial of degree at least $M+1$, irrespective of $\{y_i\}_{i=0}^{M+1}$. Denote the unique interpolating polynomial of degree $M + 1$ by $\hat{p}$. Note that the nodes of the interpolation are sorted ($y_0 < y_1 < \cdots < y_{M+1}$). Then, by the intermediate value theorem, the relation $\pm(-1)^i (f_{i+1} - f_i) > 0$ implies that a root of the derivative of the interpolating polynomial must occur for every three contiguous points. This is true since the alternating relation implies that if $f_i > f_{i-1}$, then $f_i > f_{i+1}$ as well. There are $M+2-2$ different combinations of three contiguous nodes. It follows that the interpolating polynomial has M changes of derivatives and, therefore, $\hat{p} \in \rho_{M+1}$. Hence, $\hat{p}$ can be expressed by choosing the parameters c_0, c and Y by the ansatz (4.14) as long as the nodes are sorted. Consequently, this ansatz is capable of satisfying both conditions of the theorem, finalizing the proof.

□

5 Semiclassical simulations of the laser-induced electron diffraction experiment

Obtaining accurate solutions of the TDSE is a challenging task, even for the less complex chemical reactions. This challenge is aggravated when simulations of the ever-increasing complexity of the dynamics in modern experiments are required. Nonetheless, a theoretical explanation of these experiments is required to build an understanding of the underlying physical processes.

For experiments that produce novel insights but are too complex to directly solve the TDSE for, alternative theoretical approaches are explored. The laser-induced electron diffraction (LIED) imaging technique is one of these experiments. In this chapter, we introduce the LIED experiment and use an approximate model for predicting its ouput, the semiclassical model. Instead of approximating the wavefunction at all times by finding a solution of the TDSE, the semiclassical model decomposes the dynamics of photoelectrons into classical trajectories. The statistical information of the observables is simulated by collecting the outcomes of these trajectories.

Using this approach, intrincate high-dimensional electron dynamics are approximated by solving the equations of motion for a single-active electron in the mean field of the other electrons. This creates a scalable process for complex molecules, avoiding the challenges related to their high dimensionality.

5.1 LIED imaging technique

Laser-induced electron diffraction (LIED) is a tabletop imaging technique that uses the measured photoelectron momentum distributions to determine nuclear positions with sub-bondlength precision and femtosecond timescales [134–138]. In LIED, an intense laser field, typically in the mid-infrared regime, is applied to a molecule. This laser field is selected to produce molecular dynamics in the strong field physics limit. This means that, individually, each photon does not carry enough energy to produce photoionization. Instead, electrons are ionized following other mechanisms. Under the action of the intense laser field, the molecular potential is deformed, provoking tunneling ionisation [139]. Other electrons are ionised by the action of consecutive photons, i. e., by above the threshold ionization.

After ionisation, the influece of the parent cation on the electrons is small, and the intense laser field is the driving force of the electron wavepackets. A portion of the photoelectrons change their escaping trajectory due to the laser field oscillation at each laser cycle, forcing a recollision with the parent cation. These electrons approach the nuclei of the molecule and interact with them in regions where the Coulomb force is significantly more intense than the laser field. As a result, some electrons scatter away from the molecule and some are recombined, potentially producing the ionisation of a second electron. The recombination of electrons produces the emission of coherent light in different harmonics. This process is called high-harmonic generation [140–142]. The main developers of high-harmonic generation (Pierre Agostini, Ference Krausz and Anne L'Huillier) were awarded the Nobel Prize in Physics in 2023. LIED, on the other hand, employs the electrons that do not recombine. The rescattered electrons carry information of the molecular geometry that is encoded in their escaping momentum distribution [135, 143, 144], also referred to as the photoangular distribution (PAD).

LIED stands out above other imaging techniques, such as x-ray scattering or ultrafast electron diffraction, for several reasons. First, LIED only requires an intense laser field with low-energy, and can thus be conducted in laboratories, avoiding the use of facilities that produce high-energy sources. The electrons of the molecule are used for the measurement in LIED, which improves the diferential cross section [136]. These scattering electrons in LIED have significantly lower energies than those of standard electron diffraction and are significantly perturbed by interactions with the parent ion during rescattering. This lower energetic regime results in the molecular structure not being drastically affected. High-energy reactions that destroy the molecule, such as Coulomb explosion, do not take place in LIED. The photoelectron momentum distribution encodes both information of the molecular geometry and of the molecular orbitals from where electrons were ionised [139]. Specifically, the PAD of direct photoelectrons— those that travel directly to the detector— contains information of the tunelling ionisation probability of the molecular orbitals.

The structural information of the PAD opens the possibility of gaining deeper insight into the ionization process of the molecule. But, how is it possible to separate the overlapping information from direct photoelectrons and rescaterring electrons? The separison is possible thanks to the difference in their achievable *maximum kinetic energies* [139].

A standard result in the ionisation of an electron under the action of an intense laser with peak electric field E_0 and frequency ω in the tunneling regime is that the maximum escaping kinetic energy of direct electrons is $E^{(d)}_{\text{max}} \approx 2\,U_p$ [145], where $U_p = \left(\frac{E_0}{2\omega}\right)^2$ is the ponderomotive energy. This no longer holds true for electrons that undergo scattering close to the nuclei. Similar to a rocket starting its engine in the closest point of an orbit, if the scattering occurs in the moment in which the laser field peaks and is correctly aligned, the escaping electron will drastically increase its kinetic energy. The quantitative rescatering theory (QRS) [134, 146–148] predicts that when an electron scatters with a nuclei, its outgoing momentum $\mathbf{p}_{\text{out}}$ is given by

$$\mathbf{p}_{\text{out}} \approx \mathbf{p}_r + \mathbf{A}_r, \tag{5.1}$$

where $\mathbf{p}_r$ is the rescattering momentum in the absence of a field (during rescattering), and $\mathbf{A}_r$ is the vector potential field of the laser at the time of rescattering. In the absence of a field, the molecular electric field is time-independent and, therefore, $\mathbf{p}_r$ is just a rotation of the incoming scattering momentum $\mathbf{p}_0$.

The kinetic energies of electrons are distributed in the space of spheres of radius p_r and centered at $\mathbf{A}_r$. By the QRS, the maximum incoming momentum is $p_0 = 1.26\,A_0$, where A_0 is the peak of the laser vector potential field. The maximum escaping momentum of a recolliding electron is the coherent sum of the maximum incoming momentum and the maximum vector potential, $p_{\text{max}} \sim 2.26\,A_0$. By performing the analysis with a sine wave laser field, this maximum momentum corresponds to a kinetic energy of $E^{(s)}_{\text{max}} \approx 10\,U_p$ [137, 145]. The difference in the order of the maximum kinetic energies allows us to select scattering electrons by looking at region with energies that are not accessible for direct electrons. An analysis can be applied uniquely on scattering electrons to retrieve the molecular geometry.

The duration of the laser pulses in LIED (and, therefore of electron trajectories) is in the scale of a hundred femtoseconds. The fastest vibrational frequencies in molecules are found in hydrogen compounds, such as H_2 and OH stretches, and have a magnitude of ~ 3000 cm^{-1}. This translates to vibrational periods of ~ 10 fs. An ordinary vibration frequency of molecules is around $1000\,\text{cm}^{-1}$ (tens of femtoseconds). As seen in Equation

(5.1), the electron-nucleus interactions that produce high-energy electrons take place near the peak of the laser field. In the mid-infrared, values close to the peak take place within few femtoseconds. For this reason, all high-energy photoelectron interact with a molecular geometry that is approximately identical. This explains the high spatial and temporal resolution that can be achieved by LIED.

A fundamental step towards understanding the molecular dynamics in LIED is the reconstruction of the nuclear geometry from the PAD. In the imaging community, this challenge is commonly referred to as the inverse problem. While solving the inverse problem is possible, no specific algorithm for LIED has been created to this date. Instead, inversion algorithms developed for the theory of ultrafast electron diffraction are often utilized in LIED for the similarity between imaging techniques [134, 137, 147]. This is the case of the independent atom model (IAM) [149, 150], which offers a solid approximation for retrieving the molecular geometry. The IAM approximates the molecule as a collection of atoms that do not interact with each other, neglecting the effects of correlation and molecular bonding. This approximation is valid for high-energy electrons given that the main contribution to their deflection is the influence of core electrons. Core electrons are less impacted by molecular bonding, justifying the approximation. Additionally, the IAM assumes an incoming plane wave and a single scattering event. In LIED, the recolliding wave of electrons is not necessarily a plane wave, and several scattering events are possible. For this reason, the IAM should be taken as an approximated solution to the inverse problem, and specific inversion algorithms should be developed in the future. We use the IAM for geometry retrieval in section 5.3.

To gain insight of LIED, a natural first step is to simulate the PAD from the molecular geometry. There exist quantum-mechanical simulations of the dynamics of molecular orbitals under the action of strong laser fields in the literature. However, in these simulations, the target molecules usually contain few atoms and the interest is shifted to the ionization paths, rarely containing rescattering events. Additionally, photoelectrons in LIED travel distances that are orders of magnitudes greater than the atomic scale. To resolve this, a great simulation box is required, complicating the balance between precission and computational cost. Given the complexity of solving the complete TDSE, this approach is discarded in favor of alternatives that aim to estimate the relevant observables while reducing the computational cost. The semiclassical model [138], for which we provide a description in Section 5.2, is our choice for the observable-oriented simulator.

The creation of an efficient and accurate simulation toolbox opens the path for numerical analysis of the inverse problem. By generating a number of PADs for different molecules and molecular configurations, the effect of the geometry in LIED can be interpolated using numerical techniques. Convolutional neural networks acting as encoder-decoders are one example of a possible architecture. Such numerical methods could be studied later to build an intuition on the inverse problem.

Our implementation of the semiclassical model of LIED has the potential to be used as a toolbox that motivates experimental groups for choosing the molecular samples. For this reason, we aim to create a user-friendly framework that is flexible to different choices of molecules and whose computational cost scales favorably with the system size. The bottleneck in the LIED calculations performed in this work is the electronic structure calculations, which serve as an input for the approximation of the ionisation rates and molecular electric field.

5.2 The semiclassical model

The semiclassical model is a strong-field approximation that decomposes the system dynamics into the sum of single-active-electron trajectories. The resulting observables of the molecular dynamics are estimated by the weighted sum of the individual observables of these trajectories, mimicking a histogram. Here, we describe its application to LIED. We follow the deduction in Ivanov et al. [151] and Shvetsov-Shilovski et al. [152], and refer the reader to these sources for an in-detail discussion of every step.

The dynamics of LIED are described by the solution of the electronic TDSE that has an initial condition $|\Psi_0\rangle$. This state evolves under the action of a time-dependent Hamiltonian $H(t)$, divisible into a static molecular Hamiltonian H_0 and a time-dependent perturbation $h(t)$. Let the initial condition be the ground state of the static hamiltonian $H_0|\Psi_0\rangle = E_0|\Psi_0\rangle$. The electron dynamics are resolved for fixed values of the nuclear positions, sharing the internal geometry and depending on the overall orientation. We derive the semiclassical model for a single molecular geometry first and later average out the result over nuclear orientations. For now, the geometry dependence is ommited for readability.

By the direct application of the TDSE, the wavefunction at any time is described by

$$|\Psi(t)\rangle = \int_0^t dt'\; e^{-iH(t')}|\Psi_0\rangle. \tag{5.2}$$

The perturbation is caused by an intense laser field. In the dipole approximation, the perturbation has the form $h(t) = -\mathbf{E}(t)\cdot\mathbf{x}$, where $\mathbf{E}(t)$ is the time-dependent electric field of the laser and $\mathbf{x}$ the Cartesian position operator. Using this choice of the perturbation Hamiltonian, and expanding the time-propagator in Taylor series, the solution of the TDSE can be rewritten as [151]

$$\begin{aligned}|\Psi(t)\rangle = \int_0^t dt' \left(e^{-i\int_{t'}^{t} dt'' H(t'')}\right) h(t') \left(e^{-i\int_0^{t'} dt'' H_0}\right)|\Psi_0\rangle \\ + \int_0^t dt' e^{-i\int_{t'}^{t} dt'' H_0}|\Psi_0\rangle.\end{aligned} \tag{5.3}$$

The target observable of LIED is the PAD, obtained by acting with the momentum operator $\mathbf{P}$. Let the ionic eigenfunctions of this observable with momentum $\mathbf{p}$ be denoted by $|\mathbf{p}\rangle$. The ionization condition guarantees that the state is measurable under the experimental set up. The probability of detecting an electron with momentum $\mathbf{p}$ at time t is given by the projection

$$P(\mathbf{p}, t) = \langle \mathbf{p}|\Psi(t)\rangle. \tag{5.4}$$

The detection occurs at times orders of magnitude larger than the electron dynamics timescales. For this reason $P(\mathbf{p}, t \to \infty)$ characterizes the PAD.

Given that $|\Psi_0\rangle$ is a bound state which is orthogonal to the measurable space, the projection to the second term in Equation (5.3) is zero,

$$\langle \mathbf{p}|e^{-i\int_0^{t'} dt'' H_0}|\Psi_0\rangle = e^{-i\int_0^{t'} dt'' E_0}\langle \mathbf{p}|\Psi_0\rangle = 0. \tag{5.5}$$

This term does not contribute to the PAD.

The probability of measuring an electron with momentum $\mathbf{p}$ at time t is

$$P(\mathbf{p},t) = \int_0^t dt' \langle \mathbf{p}| \left(e^{-i\int_{t'}^{t} dt'' H(t'')} \right) h(t') |\Psi_0\rangle e^{iI_p t'}, \tag{5.6}$$

where the action of the right-most operator on $e^{-i\int_0^{t'} dt'' H_0}|\Psi_0\rangle = |\Psi_0\rangle e^{i\, I_p t'}$ was substituted. Here, $I_p = -E_0$ is the ionisation potential of the molecule, i. e., the energy that is necessary for ionising the molecule.

From Equation (5.6), the conceptual functioning of the three-step model [153, 154] can be explained. The initial molecular state evolves according to the action of the static Hamiltonian H_0 until time t'. At this point, it experiences the instantaneous effect of the perbutation h, which brings the wavefunction to a different perturbed state that can be either in the continuum and/or bounded. This state then evolves under the action of the Hamiltonian until time t. By integrating over all possible interaction times (t'), all trajectories are considered.

Introduce the resolution of identity using the Volkov states of momentum $\mathbf{k}$ [155], i. e. $\mathbb{I} = \int d\mathbf{k}\ |\mathbf{k}\rangle\langle\mathbf{k}|$. The Volkov states characterize wave-like states with a certain momentum, and are constructed as eigenfunctions of the external field Hamiltonian. By introducing the identity in Equation (5.6), we obtain

$$P(\mathbf{p},t) = \int_0^t dt' \int d\mathbf{k}\ \langle \mathbf{p}| e^{-i\int_{t'}^{t} dt'' H(t'')} |\mathbf{k}\rangle\langle\mathbf{k}| h(t') |\Psi_0\rangle e^{iI_p t'}. \tag{5.7}$$

Under the strong-field low-frequency regime of LIED, the action of the perturbation Hamiltonian into the continuum states is orders of magnitude larger than the action of the molecular Hamiltonian. We introduce the approximation

$$\langle \mathbf{p}| e^{-i\int_{t'}^{t} dt'' H(t'')} |\mathbf{k}\rangle \approx \langle \mathbf{p}| e^{-i\int_{t'}^{t} dt'' h(t'')} |\mathbf{k}\rangle. \tag{5.8}$$

By definition, the Volkov states under the action of this propagator satisfy

$$e^{-i\int_{t'}^{t} dt'' h(t'')} |\mathbf{k}\rangle = e^{-i\int_{t'}^{t} dt'' E(t'')} |\mathbf{k} - \mathbf{A}(t') + \mathbf{A}(t)\rangle, \tag{5.9}$$

where $\mathbf{A}(t)$ is the vector potential of the laser field at time t and $E(t'') = \frac{1}{2}\left(\mathbf{k} - \mathbf{A}(t') + \mathbf{A}(t)\right)^2$ is the instantaneous energy at any time. Under this approximation,

$$\begin{aligned} \langle \mathbf{p}| e^{-i\int_{t'}^{t} dt'' H(t'')} |\mathbf{k}\rangle &= e^{-i\int_{t'}^{t} dt'' E(t'')} \langle \mathbf{p}|\mathbf{k} - \mathbf{A}(\mathbf{t}') + \mathbf{A}(\mathbf{t})\rangle \\ &= e^{-i\int_{t'}^{t} dt'' E(t'')} \delta(\mathbf{p} - \mathbf{k} + \mathbf{A}(t') - \mathbf{A}(t)), \end{aligned} \tag{5.10}$$

where $\delta(\mathbf{p} - \mathbf{k} + \mathbf{A}(t') - \mathbf{A}(t))$ predicts the required momentum matching of the ionised electron. Note that this condition also imposes the ionisation of the state.

Introducing this approximation in Equation (5.7), we obtain

$$P(\mathbf{p},t) = \int_0^t dt' e^{-i\int_{t'}^{t} dt'' E(t'')} \langle \mathbf{p}'| h(t') |\Psi_0\rangle e^{iI_p t'}, \tag{5.11}$$

where $\mathbf{p}' = \mathbf{p} + \mathbf{A}(t') - \mathbf{A}(t)$. The photoangular distribution is approximated by taking $t \longrightarrow \infty$ in this expression.

The semiclassical model is fully motivated from Equation (5.11). The initial state is perturbed by the action of the perturbation, providing the cumulative ionization probability given by the integral in time of $\langle \mathbf{p}'|h(t')|\Psi_0\rangle$. At some time t', an electron is ionised into a continuum state of momentum $\mathbf{p}'$. After ionization, the wavepacket evolves under the action of the Volkov state propagator. This propagator was obtained as an approximation of the complete Hamiltonian propagator, see (5.10). Reversing this approximation, the ionic wavepacket is propagated by the action of the complete Hamiltonian until the time t.

Instead of propagating Volkov wavepackets, the semiclassical model approximates these states by localised classical electrons. This step is also motivated by the path integral algorithm, in which all posible trajectories are explored and summed according to the weight of each path.

In Equation (5.11), the quantity $\langle \mathbf{p}'|h(t')|\Psi_0\rangle$ is substituted by the probability density of ionising a classical electron into the continuum at the position $\mathbf{x}'$ and with momentum $\mathbf{p}'$ at the time t', which we call $\omega(\mathbf{x}', \mathbf{p}', t')$. These electrons propagate under the action $S[\mathbf{x}]$, which now includes the intense laser field and the molecular electric field in a classical manner. We denote the classical propagator of the electrons under this action from time t' to time t as $U_H(t', t)$.

Once the action of the laser is finished, the propagator becomes time-independent and the energy of the electrons is conserved throughout the trajectories. The detected ionised electrons are those that have enough kinetic energy to escape the molecular-binding potential. For these escaping electrons, we find their asymptotic momentum $\mathbf{p}_\infty$ by using classical theory. An schematic visualisation of the three-step model of electron trajectories in LIED is shown in Figure 5.1.

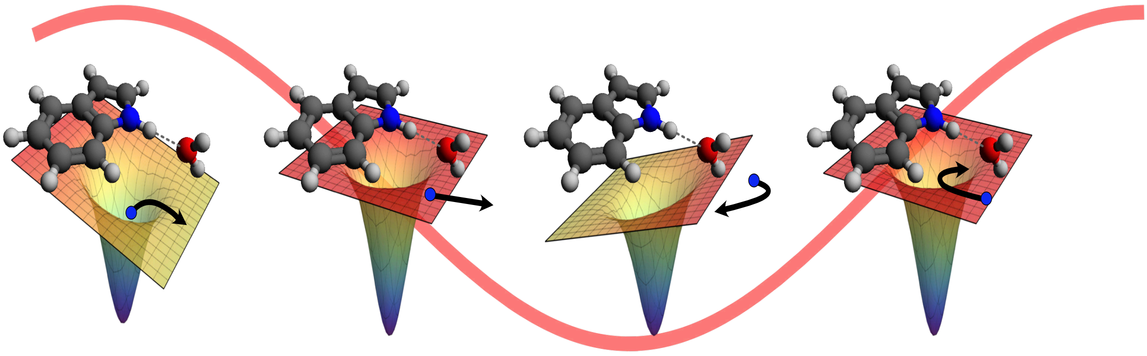

Figure 5.1: Example of a three-step trajectory of photoelectrons in LIED as a function of time, represented as the horizontal axis. Under the action of the strong laser field (depicted in red in the background), the potential of the molecule bends, inducing tunnel ionisation. The ionised electrons travel away from the molecule with the laser field as the driving force. At some point, the laser oscillates, decelerating the electrons and, in some cases, deviating them towards the molecule. These electrons travel close to the nuclei and scatter from them, possibly increasing their kinetic energy.

The approximation of the semiclassical model to the probability of measuring an electron with momentum $\mathbf{p}$ at time t is

$$P(\mathbf{p},t) = \int_0^t dt' \int d\mathbf{x}\, d\mathbf{p}' \omega(\mathbf{x},\mathbf{p}',t')\; U_H(t',t) e^{-i\int_{t'}^{t} dt'' E(t'')} e^{iI_p t'}. \tag{5.12}$$

This seven-dimensional integral is not feasible in any analytical representation. Instead, we resort to Monte Carlo integration to approximate it. For that, we create hundreds of millions initial conditions for ionised electrons according to the probability $\omega(\mathbf{x},\mathbf{p}',t')$, leading to a set of initial conditions of the photoelectron trajectories $(\mathbf{x}_0^{(i)}, \mathbf{p}'^{(i)}_0, t_0^{(i)})$, where i represents the index of the trajectory. Each electron trajectory is then solved individually according to the classical propagator $U_H(t',t)$. This is equivalent to solving the equations of motion for each trajectory.

As quantum objects, electrons carry a phase throughout their trajectories. This phase is important to reconstruct the diffraction patterns in the PAD. This classical phase mimics the exponentials of Equation (5.12). Following Shvetsov-Shilovski et al. [152], the classical phase of each electron trajectory can be approximated by

$$\Phi(t) = I_p(t_0)\, t_0 - \mathbf{p}_0' \cdot \mathbf{x}_0 - \int_{t_0}^{t} dt'' \left(\frac{\mathbf{p}^2(t'')}{2} - \frac{2Z}{r(t'')} \right), \tag{5.13}$$

where $I_p(t_0)$ is the field-dependent ionisation potential at time t_0, Z is the cationic charge of the molecule after the ionisation event (usually $Z = 1$), and r is the radial distance to the origin of coordinates located at the centre of charge $r = \|\mathbf{x} - \mathbf{x}_e\|$. We define the initial phase as $\Phi_0 = I_p(t_0)\, t_0 - \mathbf{p}_0' \cdot \mathbf{x}_0$.

As a last step, we average out these results over the possible nuclear rotational configurations. To account for this degree of freedom, we include the angular dependency on the ionisation rates. The modified version of Equation (5.12) that includes all possible laboratory frame rotations is

$$P(\mathbf{p},t) = \int_0^t dt' \int d\mathbf{x}\, d\mathbf{p}'\, d\Omega\, \omega(\mathbf{x},\mathbf{p}',t',\Omega)\; U_H(t',t;\Omega)\; e^{-i\int_{t'}^{t} dt'' E(t'')} e^{iI_p t'}, \tag{5.14}$$

where Ω is the set of the three Euler angles that parametrise the overall rotations of the molecule. As a result of choosing an orientation-dependent ionisation theory, the propagator $U_H(t',t;\Omega)$ becomes orientation-dependent.

Overall rotations can produce alignments of the molecular orbitals with the electric field that facilitate or impede ionisation, thus producing higher ionisation yield for some orientations. In some experimental setups, the molecular sample is aligned in the direction of some electric field prior to the ionisation step. This effect is completely independent of the ionisation rate of a molecule and Bayesian probability is required. The ionisation rate of an electron with the nuclear orientation given by Ω under the action of aligning field that induces an alignment probability distribution $\rho_{\text{align}}(\Omega)$ is

$$\omega(\mathbf{x},\mathbf{p}',t',\Omega) = \rho_{\text{ion}}(\mathbf{x},\mathbf{p}',t',\Omega)\; \rho_{\text{align}}(\Omega), \tag{5.15}$$

where ρ_{ion} refers to the ionisation probability due to the laser field at the orientation Ω.

Equation (5.14) is a ten-dimensional integral that can be approximated by using Monte Carlo algorithms with a great number of sampled electrons, as described in subsection 1.5.2. For generating the initial electron conditions, we use rejection sampling. To acknowledge for the volume element in the angular coordinates, a multiplication by $\sin(\beta)$ is added to the integrand, where β is the rotation around the y axis following to the zyz convention.

5.3 Results

5.3.1 Computational setup

We provide an example of the results of a complete simulation of the PAD of the LIED experiment. The molecule in this example is OCS, resting in its equilibrium geometry, given by the three-dimensional position of the atoms $\mathbf{x}_\mathrm{O} = [0.0, 0.0, 3.179]$, $\mathbf{x}_\mathrm{C} = [0.0, 0.0, 0.989]$ and $\mathbf{x}_\mathrm{S} = [0.0, 0.0, -1.961]$ in atomic units. In this configuration, the ionisation potential of the molecule is $I_p = 11.2$ eV $= 0.4112$ au [156]. The non-zero components of the permanent dipole and polarizability of the neutral OCS in this configuration are [157] $\mu_z = -0.72$ D, and $\alpha_{x,x} = 26.15$ au, $\alpha_{y,y} = 26.15$ au and $\alpha_{z,z} = 50.72$ au. For the cationic OCS, the non-zero components of the permanent dipole and polarizability are [157] $\mu_z = -2.22$ D, and $\alpha_{x,x} = 19.06$ au, $\alpha_{y,y} = 18.73$ au and $\alpha_{z,z} = 44.09$ au.

A linearly-polarised laser of intensity $I = 1 \cdot 10^{14}$ W cm^{-2} was used. This intensity is equivalent to a peak electric field of $E_0 \approx 0.0534$ au. The wavelength of the laser was 2000 nm, which translates to an angular frequency of $\omega \sim 0.02278$ au. The laser acts for a total of 11 cycles, which translates into a total duration $T \approx 73.3841$ fs ~ 3033.7995 au. This is equivalent to a full-width half maximum duration of roughly 37 fs. These values were chosen to resemble the ionizing electric field of Karamatskos et al. [137]. Under this choice, the ponderomotive energy is $U_p \sim 37.3759$ eV and the kinetic energy of returning electrons (roughly 3.17 U_p) is ~ 118.4466 eV. The laser field was analytically approximated in accordance to (5.21).

The molecular sample was not aligned prior for the calculation. In other words, the molecular sample was simulated as isotropically distributed.

The ionisation probabilities were computed by combining the molecular-orbital applied Ammosov-Delone-Krainov (MO-ADK) theory, as described in subsection 5.5.2, and the transverse distribution of the molecular orbital, described in subsection 5.5.3.

The molecular electric field for this molecular geometry was obtained using the ESP package, as described in subsection 5.5.4.

The propagation of the semiclassical trajectories were performed by solving the equations of motion (EOMs) introduced in subsection 5.5.5. Concretely, the EOMs are described in Equation (5.44). Given that the EOMs form a system of coupled PDEs that can be stiff, a 5-th order implicit Runge-Kutta propagator [158] was used.

After propagation, the detection of electron trajectories by a velocity-map imaging spectrometer was simulated as described in subsection 5.5.6. From the detection step, the PADs were obtained.

Given that the trajectories of the semiclassical model are independent to each other, this algorithm can be naively parallelised over different CPUs / nodes. Each node was used to sample and propagate 10^5 electrons at a time. Each job took roughly 10 minutes, that is 6 ms per electron trajectory. A total of 10^9 electrons were propagated for this simulation. Using 200 non-exclusive nodes, the complete semiclassical model of LIED took approximately 10 hours. As the LIED package allows for flexible choices of the calculation of the ionisation rate and molecular electric field, we consider these calculations as an input to the LIED model and did not include them in the reported running time.

5.3.2 Results and geometry retrieval

The simulated PAD of LIED for the isotropically distributed OCS molecular sample in its equilibrium geometry is shown in Figure 5.2(a). The section of the PAD containing electrons with maximum kinetic energy after rescattering ($10\ U_p$) are marked with a blue dotted circle. A preferential direction of the high-energy yield in the laser polarization axis (the z axis) can be observed.

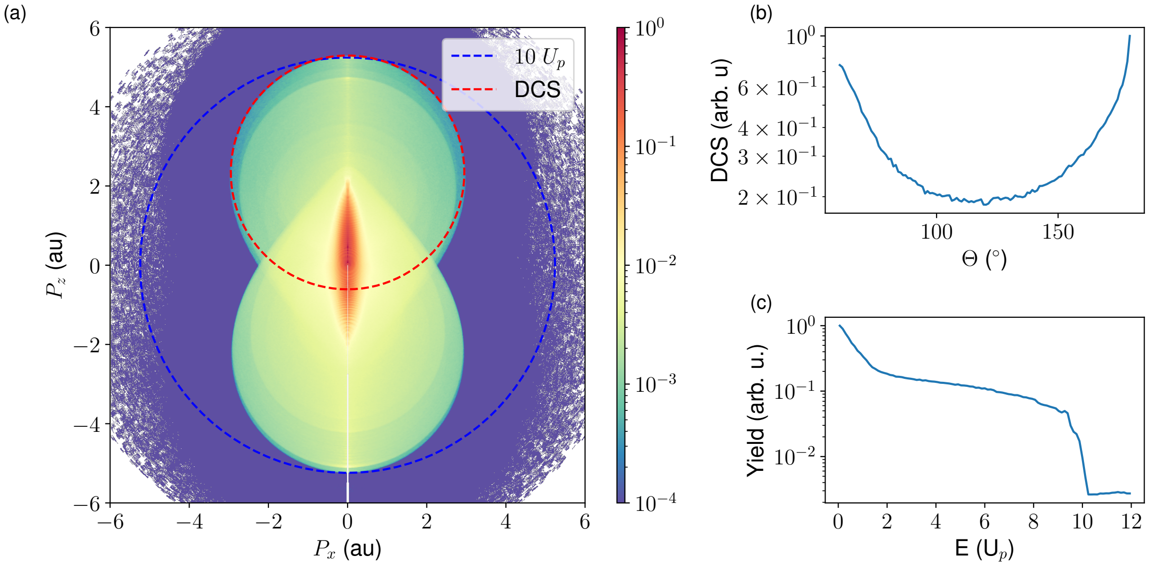

Figure 5.2: **(a)** Simulated photoangular distribution of OCS at its equilibrium geometry (arbitrary units). The PAD was obtained by propagating the trajectories of 10^9 photoelectrons. The simulated image was normalized such that the maximum intensity is 1. Two circles were added to the PAD to show the position of electrons with $10\ U_p$ kinetic energy (blue), and of electrons that are used in the differential cross section (DCS) of the independent atom model (red). **(b)** Electron yield of electrons that have rescattered near the peak of the vector potential field of the laser. The DCS was obtained by performing a histogram of the trajectories around the momentum defined in (5.16) using an integration box of $\Delta p = 0.05\ p_0$ and $\Delta\theta_r = 1°$. **(c)** Electron yield with respect to the kinetic energy of photoelectrons.

By the quantitative rescattering theory (QRS) [134, 146–148], returning electrons can have incoming momentum at most $p_0 = 1.26\ A_0$. If an electron returns with maximum momentum and rescatters at the peak of the vector potential field, its escaping momentum is given by

$$\begin{cases} p_x = -1.26\ A_0 \sin(\theta_r), \\ p_z = -1.26\ A_0 \cos(\theta_r) + A_0, \end{cases} \tag{5.16}$$

where θ_r is the rescattering angle. Photoelectrons achieve their maximum escaping kinetic energy at backscattering ($\theta_r = \pi$). The electron yield at this momentum is used to obtain the differential cross section of the LIED experiment [159]. The position of the DCS electrons is marked in Figure 5.2(a) by a red dotted circle. The figure displays two such rings, corresponding to interactions of electrons with maximal returning energy and opposite directions. These rings accurately match their predicted position.

The electron yield of the PAD in the DCS momentum region for OCS is shown in Figure 5.2(b). To obtain the DCS, the PAD was evaluated at an integration box of $\Delta p = 0.05\, p_0$ and $\Delta\theta_r = 1^\circ$ centered around the momentum values in (5.16).

Additionally, the yield with respect to the kinetic energy of photoelectrons is shown in Figure 5.2(c). The results show the expected tendency [146, 160]: the photoelectron yield decays quickly for electron with kinetic energy until roughly $3\, U_p$. Then, a plateau is found for kinetic energies between 4 to $10\, U_p$. In this region, the electron yield becomes almost constant. Last, a sharp decay is once more seen in the region with energies larger than $10\, U_p$, pointing out that this is the limit for the kinetic energy of escaping electrons.

An overestimation of the ratio between the yield of electrons with energy $E > 3U_p$ to those of $E < 3U_p$, compared to experimental results [137], can be observed. This can be associated to two reasons. First, the simulation was performed for the laser field corresponding to the peak of the spatial profile. In other words, the spatial distribution of the laser was assumed to be a delta function on the peak, instead of the usual Gaussian distribution. To account for the spatial profile, subsequent simulations with smaller peak electric fields can be performed. Simulations with decreasing peak laser fields generate electrons with decreasing maximum energy. Therefore, this combination generates an electron yield with a smaller ratio of high-energy electrons. The second source of error is the approximations performed for the momentum distribution of tunneling electrons, see subsection 5.5.3 for a discussion.

Electrons can only access the DCS momentum in (5.16) if they interact with the laser at the peak field. Therefore, the lack of spatial profile integration does not impact the relative yield of electrons in the DCS. The structure retrieval analysis can be performed on these electrons.

We exemplify the application of the independent atom model (IAM) [134, 161] to obtain the molecular geometry from the DCS. For its application in LIED, an initial incoming planewave is assumed. This is consistent with the quantitative rescattering theory that approximates the fastest returning electrons as having momentum $(p_0)_z = -1.26\, A_0$. Under these assumption, the IAM approximates the DCS at momentum $\vec{p}$ for a molecule composed of N atoms as [134, 137, 147]

$$I(\vec{p}, \theta_r) = W(\vec{p}_0) \sum_{i,j=0}^{N} \int d\Omega\, P(\Omega)\; f_i(\theta_r) f_j^*(\theta_r)\; e^{i\vec{q}\vec{R}_{ij}(\Omega)}, \tag{5.17}$$

where $P(\Omega)$ is the angle-dependent ionization probability, $f_i(\theta)$ is the atomic contrast factor of the i-th atom, $\vec{q} = \vec{p} - \vec{p}_0$ is the momentum transfer, and $\vec{R}_{ij} = \vec{R}_i - \vec{R}_j$ is the matrix of distances between atoms. The distance matrix is dependent of the overall orientation of the molecule Ω.

$W(\vec{p}_0)$ is a function that encodes the momentum distribution of the incoming electron packet. To get rid of the dependence to this distribution, the molecular contrast factor (MCF) is introduced. For that, the contribution of the isolated atoms to the differential cross section (I_{atom}) is defined by the sum of the $i = j$ terms in the DCS, that is

$$I_{\text{atom}}(\theta_r) = W(\vec{p}_0) \sum_{i=0}^{N} |f_i(\theta_r)|^2 \int d\Omega\, P(\Omega). \tag{5.18}$$

The atomic contribution does not depend on the outgoing momentum or position of the nuclei.

Using I_{atom}, the MCF is defined as the ratio of variation of the DCS from the atomic DCS by the interference between different atoms. The MCF is given by [134]

$$\gamma_{\text{MCF}} = \frac{I - I_{\text{atom}}}{I_{\text{atom}}} = \frac{\sum_{i \neq j}^{N} \int d\Omega \, P(\Omega) \, f_i(\theta_r) f_j^*(\theta_r) \, e^{i\vec{q}\vec{R}_{ij}(\Omega)}}{\sum_{i=0}^{N} |f_i(\theta_r)|^2 \int d\Omega \, P(\Omega)}. \tag{5.19}$$

In this step, the necessity of a knowledge of the distribution $W(\vec{p}_0)$ was removed.

The molecular geometry is fitted by minimizing the distance between the MCF obtained from the LIED simulation ($\gamma_{\text{MCF}}^{(S)}$) and the theoretical MCF ($\gamma_{\text{MCF}}^{(T)}(\vec{R}_i)$). In other words, the molecular geometry is retrieved by optimizing the distance function $\chi^2(\vec{R}_i)$ given by

$$\chi^2(\vec{R}_i) = \left(\gamma_{\text{MCF}}^{(S)} - \gamma_{\text{MCF}}^{(T)}(\vec{R}_i)\right)^2 = \left(\frac{\beta \, I^{(S)} - I^{(T)}(\vec{R}_i)}{I_{\text{atom}}}\right)^2, \tag{5.20}$$

where $I^{(S)}$ is the DCS obtained from the LIED simulations and β is a normalization constant. The normalization constant allows for the scaling of the histogram containing the DCS to theoretically acceptable values. A different normalization constant is used for each configuration during the fitting. If the MCF is computed on a grid of G angles $\{(\theta_r)_g\}_{g=1}^{G}$ with a given importance weight, β can be analytically calculated. For our calculations, the scattering angles are equidistant, satisfy $\theta_r > 60°$, and the weight is constant throughout the different angles.

The atomic form factors for the calculations were obtained using the ELSEPA package [162]. These form factors were computed using an incoming electron wavepacket at 100 eV, which quantitavely matches the kinetic energy of returning electrons. To account for the molecule being in its cationic states when the rescattering takes place, the form factor of the sulfur atom was computed in its cationic state S^+. The form factor of the other atoms were calculated in their neutral form. This charge distribution provides the best results [137] as sulfur presents the lowest ionization energy of the atoms in OCS.

An ionization theory is needed for the integration of the angular dependency of the MCF in Equation (5.19). Here, we use the experimental ionization probabilites reported in Karamatskos et al. [137]. These experimental ionisation probabilities depend on the angle between the laser polarization axis and the aligning field axis. The effect of the other Euler angles were averaged out in the experimental measurement step.

The results of the application of the IAM atom model to the DCS of OCS using experimentally reported ionization rates and MO-ADK theory is presented in Figure 5.3. A linear configuration was assumed for the calculation and the χ^2 measure was computed using different bond distances. The geometry that was inputed to the calculationd and the geometry with the best fitting error are marked in the figure with a black dot and white cross, respectively. The obtained error for the O-C bond length is 0.1 pm, and for the C-S bond length is 9.5 pm.

5.4 Conclusions

In this chapter, we succesfully implemented the semiclassical model of LIED. We showcase the capability of the experiment to predict molecular geometries with high spatial and temporal resolution. This is done by performing a simulation for a given geometry and

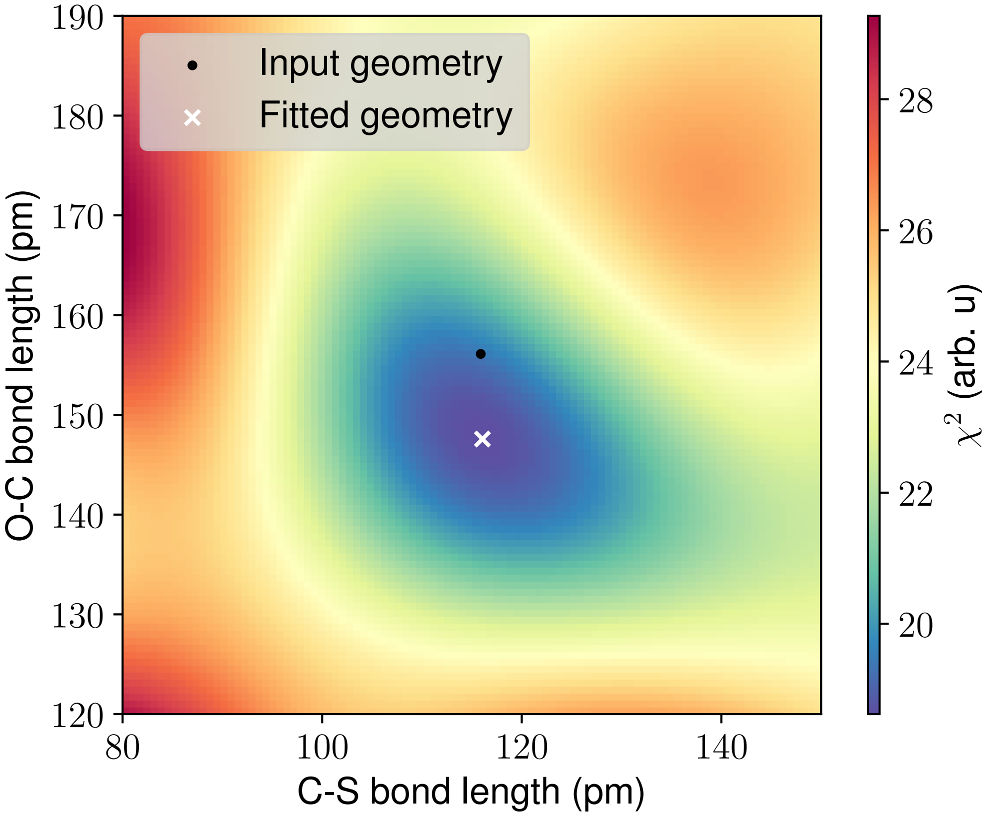

Figure 5.3: χ^2 error as a function of the bond-lengths of the OCS molecule. A linear configuration was assumed for the calculation. The input geometry of the calculation is marked with a black dot and the geometry with the smallest error is marked with a white cross. The distance between results is on the order of 10 pm.

analysing its result with the IAM. The implemented model is flexible to different choices of molecular sample, molecular orbital, alignment state and laser field. The impact of ionisation from different molecular orbitals can be averaged by performing subsequent calculations.

While the current semiclassical model is capable of reproducing a DCS that retrieves the geometry of the molecule, two main improvements would help it match the experimental results more accurately. First, the laser spatial distribution can be taken into account by integrating over the results of LIED with different peak laser field intensities. Then, the non-adiabatic ionization theory used for the initial momentum distribution of photoelectrons can be substituted by a molecular-orbital dependent theory. Both of this additions would improve the ratio of high-energy electrons, reproducing the experimental results with higher accuracy.

5.5 Appendix: Implementations of the required functions of LIED

Several different pieces are needed to assemble the semiclassical model. This appendix serves for a description of these individual pieces. The semiclassical model requires the creation, propagation and detection of hundreds of millions of electrons. Therefore, the time and memory efficiency of each individual piece is fundamental. Improvements in computational efficiency are often accompanied by a trade-off in accuracy. The semiclassical model is assembled using algorithms that provide a good accuracy-to-efficiency ratio.

First, the experimental setup is defined. The experimental quantities are the ionising laser field and the alignment probability distribution. The analytical approximation of the experimental quantities is discussed in subsection 5.5.1.

The initial conditions for the electron trajectories require a ionization theory in intense laser fields to be created. This theory needs to allow for quick evaluations at different orientations. The molecular-orbital applied Ammosov-Delone-Krainov (MO-ADK) theory was chosen. MO-ADK is described in subsection 5.5.2. An additional theory is required to predict the position and momentum of the electron at ionization. This theory is discussed in subsection 5.5.3.

To propagate electron trajectories, the electric field of the molecule is required. In order to reduce the complexity of the approximation, we narrow the focus to describing the electrostatic potential (ESP) field. Having approximated the ESP, the electric field of the molecule can be found as its gradient over space. Thanks to automatic differentiation techniques, this gradient can be numerically computed from the analytical representation of the ESP. We describe different approaches to approximate the ESP and provide an analysis of their accuracy in subsection 5.5.4. Using the electric field, the electron trajectories are propagated. The setup of the algorithms used for propagation is detailed in subsection 5.5.5.

The electrons that escaped the molecular binding potential continue their asymptotic movement until they are detected. For detection, a velocity-map imaging (VMI) detector is used. The simulation of the detection process is summarised in subsection 5.5.6.

5.5.1 Experimental quantities: laser field and alignment

An important part of simulating any experiment is replicating its setup in the form of an analytical digital twin. Such a representation is usually unavailable and simplified approximations that depend on few parameters are commonly used.

In this work, we assume that the ionising laser field is linearly polarized and choose to align it with the z axis. In theory, laser fields have an infinite duration and their intensity decays according to some exponential rule. This is not a practical choice as it requires infinitely-long timespans to account for the complete action of the laser. For this reason, we choose a cosine-squared envelope that creates a zero valued laser-field at finite time. The laser field with total duration T, pulse frequency ω and peak electric field E_0 centered around $t = 0$ is given by

$$E_z(t) = \begin{cases} E_0 \cos(\omega t) \cos^2\left(\frac{2\pi t}{T}\right), & \text{if } |t| \leq T/2 \\ 0, & \text{else.} \end{cases} \tag{5.21}$$

Alternatively, the laser can be defined over a number of cycles n, following the relation $n = T\omega/(2\pi)$.

Having an analytical expression for the electric field simplifies the calculation of the vector potential field, which is needed in the LIED calculations. The vector potential field is also aligned in the z axis and is minus the antiderivative with respect to time of the electric field.

The electric field and vector potential field of a linearly-polarised laser under a cosine-squared envelope can be seen in Figure 5.4. The created laser is symmetric from its center and, the vector potential is antisymmetric. The figure illustrates that the vector

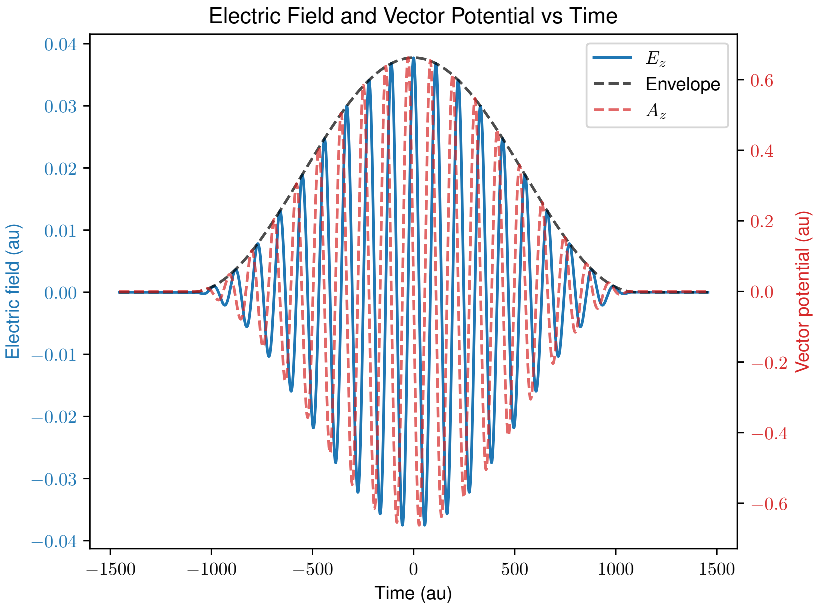

Figure 5.4: Z-component of the electric field and vector potential field of a laser generated with a cosine-squared envelope as in Equation (5.21). The laser has frequency $\omega = 800$ nm, intensity $\mathrm{I} = 5 \cdot 10^{13}\ \mathrm{W/cm^2}$, equivalent to $E_0 = 0.0377$ au, and a duration of 20 cycles, equivalent to $T = 2206.4$ au.

potential is a quarter of a cycle delayed from the electric field, which is consistent with the common plane-wave behaviour. The decay into a value that is exactly zero at the limits can also be seen on the figure, guaranteeing that the laser has a finite time-span.

As part of the experimental preparation of the molecular sample, the molecules can be aligned according to some external field. While the dynamics of the molecular alignment are beyond the scope of this work, the degree of alignment affects the obtained photoangular distribution. Given a molecular sample that was aligned in the direction $\mathbf{u}$ and has a degree of alignment $\langle \cos^2 \theta \rangle$, the alignment probability with respect to the orientation is given by

$$\rho_{\text{align}}(\Omega) = \exp\left(-\frac{\sin^2(\angle \mathbf{ud}(\Omega))}{2(1 - \langle \cos^2 \theta \rangle)}\right), \tag{5.22}$$

where $\mathbf{d}(\Omega)$ is the vector that corresponds to the rotation of the molecular dipole by the Euler angles Ω, and $\angle \mathbf{ud}(\Omega)$ is the angle between $\mathbf{u}$ and $\mathbf{d}(\Omega)$.

If the molecular sample is not aligned and follows an isotropic distribution, then $\rho_{\text{align}}(\Omega)$ has a constant value.

5.5.2 Ionisation probability: MO-ADK

Computing the ionisation rate of molecules is necessary for the generation of initial conditions for electrons. For the creation of hundreds of millions of trajectories, we require a theory that allows for an efficient evaluation. In addition, the theory needs to allow for different orientations of the molecule. MO-ADK satisfies all of these requirements, as it relies on a single fitting step of coefficients that can then be stored and reused for subsequent calculations. Here, we provide a brief description of the model and refer the reader to Zhao et al. [163] and Tong and Lin [164] for a deeper insight.

The electronic wavefunction of the single-active electron in the single-centre expansion approach can be divided into the radial and angular parts as

$$\Psi(\mathbf{r}) \approx \sum_{l,m,n} C_{lmn} R_n(r) Y_{lm}(\theta, \varphi), \tag{5.23}$$

where r is the distance to the centre of electric charge, θ is the polar angle and φ is the azimuthal angle. The functions Y_{lm} are the spherical harmonics and C_{lmn} are the structure parameters.

In the asymptotic region, the radial part approaches a single function that takes the form

$$R(r \longrightarrow \infty) \sim \; r^{Z_c/\kappa - 1} e^{-\kappa r}, \tag{5.24}$$

where Z_c is the effective charge of the molecule and $\kappa = \sqrt{2I_p}$. In the case of neutral molecules, the electron in the asymptotic region perceives the molecule as a cation, therefore $Z_c \approx 1$. As on the assymptotic region only one radial function is required, the n quantum number is omitted from this point on.

The values of the single-active electron orbital at the asymptotic region can be used to fit the effective charge and structure parameters. Our implementation of MO-ADK fits Z_c using gradient-descent methods and finds the best set of coefficients C_{lm}, up to order l_{max}, by solving an overdetermined set of linear equations. To obtain the values of the wavefunction at different points, we include the framework for setting up quantum-chemistry calculations using the Psi4 package [100]. This implementation serves as a possible choice for generating the input to MO-ADK. However, other software can be used to generate this input. This enables a high flexibility on the choice of quantum-chemistry theory, allowing for the possible utilization of user-built and/or molecule-specific calculations.

For a given fit of effective charge and set of structure parameters, the ionisation rate of a molecule under the influence of a strong laser field $\mathbf{E}$ along the molecular axis, can be computed using MO-ADK theory as

$$\omega_{\text{field}}(\mathbf{E}, \Omega) = \sum_m \frac{|B(m,\Omega)|^2}{2^{|m|}\,|m|!} \frac{1}{\kappa^{2Z_c/\kappa - 1}} \left(\frac{2\kappa^3}{E}\right)^{2Z_c/\kappa - |m| - 1} e^{-2\kappa^3/3E}, \tag{5.25}$$

where E is the norm of the electric field and

$$B(m, \Omega) = \sum_l (-1)^{(m+|m|)/2}\, C_{lm}(\Omega) \sqrt{\frac{(2l+1)(l+|m|)!}{2(l-|m|)!}}. \tag{5.26}$$

The definition of $B(m, \Omega)$ is derived from the approximation of the spherical harmonics by the leading term of the expansion.

Given the use of spherical harmonics, the created quantities can be computed for different orientations by rotating the coefficient set $C_{lm}(\Omega)$. For expressing the rotation in the coefficients, the Wigner matrices $W^l_{m,m'}(\Omega)$ are used, where Ω are the Euler angles associated to the rotation. The rotated coefficients are given by

$$C_{lm}(\Omega) = \sum_{m'} C^{(0)}_{lm'} W^l_{m,m'}(\Omega), \tag{5.27}$$

where $C^{(0)}_{lm'}$ are the fitted coefficients, originally calculated for a given state corresponding to $\Omega = 0$.

The cumulative probability of ionisation of an electron under the action of a laser pulse can be computed by integrating the ionisation rate (using the time-dependency of the laser field) over time as

$$P(t, \Omega) = 1 - \exp\left(-\int_{t_0}^{t} dt' \; \omega_{\text{field}}(\mathbf{E}(t), \Omega)\right). \tag{5.28}$$

The probability density of ionization ρ_{field} can be computed as the time derivative of this quantity. Its expression is

$$\rho_{\text{field}}(t, \Omega) = \omega_{\text{field}}(\mathbf{E}(t), \Omega) \cdot \exp\left(-\int_{t_0}^{t} dt' \; \omega_{\text{field}}(\mathbf{E}(t), \Omega)\right). \tag{5.29}$$

This equation shows that, to compute the probability of ionisation, the ionisation rates at all times prior to t are required for integration. This hinders the efficiency of the method, as sequential evaluations of ω_{field} need to be computed for using some integration scheme. To improve the efficiency of this calculation in the LIED simulations, we use the approximation $\rho_{\text{field}}(t, \Omega) \approx \omega_{\text{field}}(\mathbf{E}(t), \Omega)$.

The application of the MO-ADK theory is illustrated for the water molecule. An ionising laser field with the same frequency, intensity and duration as the depicted in Figure 5.4 is applied to the molecule. The expansion coefficients for different quantum-chemistry methods and different basis set sizes were fitted to the values of the highest-occupied molecular orbital (HOMO) in order to visualize the convergence of the results. The total ionisation probability is computed by integrating the action of the laser and obtained with respect to the angle between the laser and the water dipole. The other angles were averaged out. The results for the MO-ADK ionisation probability are shown in Figure 5.5. It can be seen that the different choices of method and basis set sizes produce similar results, pointing out that the calculations are convergent. We compare the fitted results from our calculations to those obtained with the previously reported coefficients in Zhao et al. [163]. The fitted ionisation rate produces similar results.

MO-ADK theory offers an analytical expression for the ionisation rate that can be applied to any molecule under strong field lasers. However, it must be noted that several problems are found when using it in practical calculations. The focus of MO-ADK is the behaviour of the wavefunction in the asymptotic region, i. e., at long ranges where the wavefunction approaches zero. This radically oposes the traditional quantum-chemistry approaches, centered on finding the molecular orbital on small distances to nuclei, where the molecular orbital density is greater. Calculations are performed using sums of atomic orbitals centered at the nuclei, whose radial part decay as $\sim e^{-r^2}$. The effect of the wavefunction decay at long ranges is minimal for the convergence of electronic energies, provoking that the right asymptotic behaviour is, in practice, never achieved. This harshly impacts the stability of the fitting of the effective charge and structure parameters in MO-ADK. This effect was manifested when fitting for different minimum distances to the molecular center.

MO-ADK is only capable of estimating the ionisation rate for a single molecular orbital. Using the HOMO as the only ionizable orbital is, in general, an acceptable approximation. However, as the ionizing laser field becomes more intense, the contribution

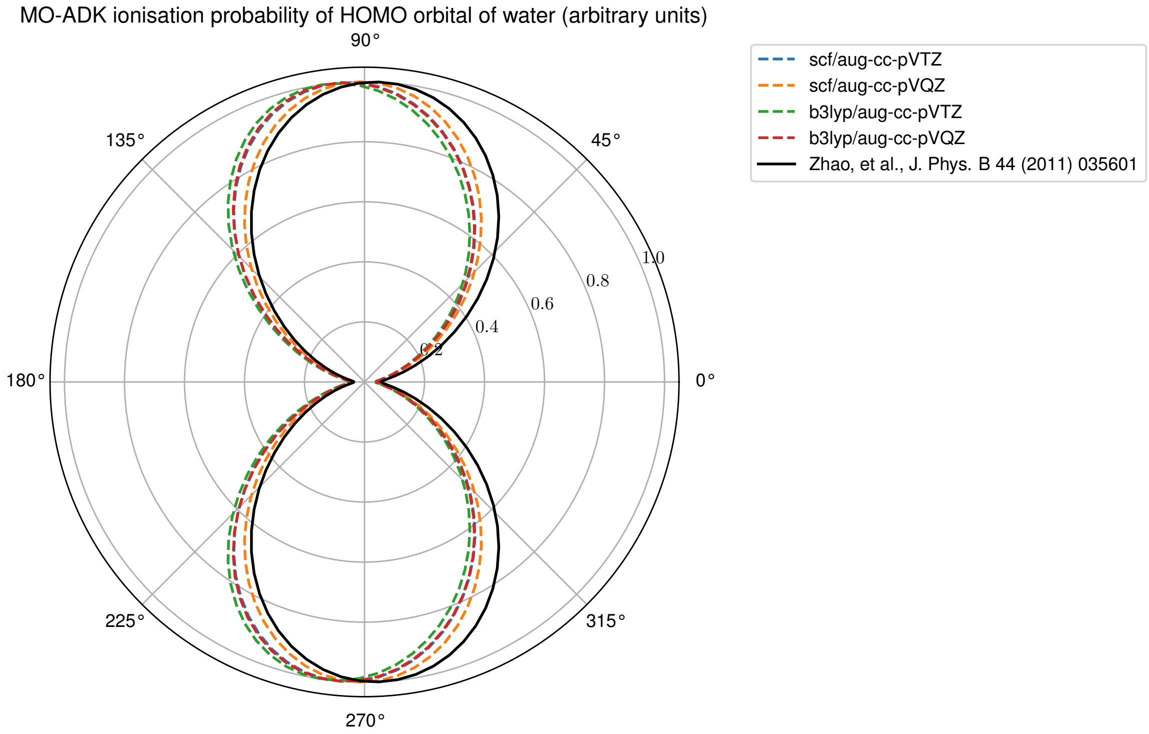

Figure 5.5: Ionisation probability of the HOMO orbital of the water molecule with respect to the angle between the laser field and the dipole of the molecule. The other Euler angles were averaged out. Results were computed using the MO-ADK theory under the action of a laser field of frequency $\omega = 800$ nm, intensity $\mathrm{I} = 5 \cdot 10^{13}$ W/cm^2 and duration of 20 cycles. The calculations performed for the self-consistent field (scf) and the Becke 3-parameter Lee–Yang–Parr (b3lyp) methods are shown for two choices of basis set size (aug-cc-pVTZ and aug-cc-pVQZ). Additionally, the results using the expansion coefficients of water reported in Zhao et al. [163] are displayed for comparison. All the result are created from fitting coefficients up to $l_{max} = 5$.

of other orbitals to ionization becomes larger. This problem can be overcome by performing a sum of the ionization rates of the different orbitals that yield ionisation. Another alternative is using the Dyson orbital for the fitting, i. e., the overlap between the N-electrons and the $N-1$-electrons wavefunctions.

5.5.3 Initial conditions of photoelectrons

MO-ADK provides the instantaneous ionization probability of a molecule under the action of an intense laser field. However, no information is provided about the dependence of the probability on the position and momentum of the electrons. To find this dependence, tunneling ionization theory is required. We use a simplified version of the theory in Murray et al. [165] for the probability of finding a photoelectron with position $\mathbf{x}$ and momentum $\mathbf{p}$.

Due to the alignment of the laser on the z direction, the exiting position z_{exit} of the photoelectron in the tunnel ionization is centered around

$$z_{\text{exit}} = -I_p/E_z, \tag{5.30}$$

where E_z is the laser field at the moment of ionization. Small deviations from the exiting position are possible, but, in practice, their contribution to the dynamics is neglectable.

The tunneling momentum of electrons requires a more complex analysis. Let $\Psi(x, y, z)$ denote the Dyson orbital of the molecule. Then, denote by z_0 the matching point of the tunnel ionization, i. e., the point where the potential barrier is maximal. Evaluate Ψ at z_0 and perform the Fourier transform for the perpendicular coordinates to obtain

$$\Phi(p_x, p_y, z_0) = \frac{1}{2\pi} \int dx\, dy\; e^{-ixp_x - iyp_y} \Psi(x, y, z_0) \tag{5.31}$$

By performing a WKB propagation of the Fourier transformed Dyson orbital to the continuum, the Dyson orbital in the space coordinates yields

$$\Psi(x, y, z) \approx \frac{1}{2\pi} \sqrt{\frac{\kappa}{|p_z(z)|}} e^{-\int_{z_0}^{z} dz'\, p_z(z')} \int dp_x\, dp_y\; e^{-ixp_x - iyp_y} \Phi(p_x, p_y, z_0) e^{-\frac{1}{2}(p_x^2 + p_y^2)\tau}, \tag{5.32}$$

where $\tau = \kappa/E$.

The tunneling ionization rate depends on the absolute squared value of the Dyson orbital at the tunneling exit point. The ionization rate is given by[4]

$$\rho_{\text{tunnel}}(\mathbf{x}, \mathbf{p}', t') \sim e^{-2\int_{z_0}^{z_{\text{exit}}} dz'\, p_z(z')} \left|\Phi(p_x, p_y, z_0)\right|^2 e^{-(p_x^2 + p_y^2)\tau}\; \delta(\mathbf{x} - I_p(t')/\mathbf{E}(t')). \tag{5.33}$$

The probability density in the perpendicular momentum components is governed by a Gaussian of variance $1/\tau$, which cuts the contribution for high-momentum electrons. In addition, the dependence on $\Phi(p_x, p_y, z_0)$ shows that the perpendicular momentum is benefitted from the Dyson orbital being extense in the direction of the electric field.

The rate depends on the parallel momentum via $e^{-2\int_{z_0}^{z_{\text{exit}}} dz' p_z(z')}$, which also incorporates the dependency with the shape of the barrier. The evaluation of this exponential can be avoided by searching for alternative theories. A non-adiabatic approximation for the parallel momentum at tunneling is [166]

$$p_z = -A_z \left(\sqrt{1 + \gamma^2} - 1\right), \tag{5.34}$$

where A_z is the vector potential of the laser at tunneling and $\gamma = \frac{\omega}{|E_z|}\sqrt{2I_p + p_x^2 + p_y^2}$ is the effective Keldysh parameter. Then, longitudinal momentum is determined in terms of the other momentum components and the tunneling time.

Even in this non-adiabatic approximation, the ionization rate depends on the matching point, which is a time-dependent quantity (as it depends on the laser field). Finding this point requires iterative methods even for static fields. To avoid this, we use an approximate arbitrary choice of this point. We assume that the differences in this value does not lead to significant changes in the theory. This approximation holds true as long as the coherent sum of spherical harmonic that represents the Dyson orbital does

---

[4] At the moment of writing this dissertation, the dependence with respect to the Fourier transform of the Dyson orbital has not yet been implemented in the LIED software. Instead, the approximation $\left|\Phi(p_x, p_y, z_0)\right|^2 = 1$ was used. This approximation does not affect the possible outcomes of electron trajectories, but it affects the relative density. The implementation of this dependency is planned as the first upgrade to the current software.

not depend on this point. We define $z_0 = z_{\text{exit}} - \Delta z$, where Δz is an arbitrary constant with a value roughly half the barrier width.

With this, the recipe for the total ionization rate depending on the ionising laser field, the photoelectron momentum and position is complete. The total ionisation rate is given by

$$\rho_{\text{ion}}(\mathbf{x}, \mathbf{p}', t', \Omega) = \rho_{\text{field}}(\mathbf{E}(t'), \Omega)\, \rho_{\text{tunnel}}(\mathbf{x}, \mathbf{p}', t'), \tag{5.35}$$

where the respective field and tunneling ionization rates are defined in Equations (5.25) and (5.33).

5.5.4 Electrostatic potential: ESP

While the laser field is the main driving force of the electron trajectories on the long-distance regime, the molecular electric force takes a crucial role for incoming electrons at distances smaller than few angstroms from the nuclei. The nuclei are considered static, making the nuclear-electron interaction analytically expressable using Coulomb's law. This term will be dominant for electron rescattering. The electron-electron interaction with the parent cation is important for the correct characterization of the photoelectron trajectories but no analytical expression is available. In this subsection we explore the fitting of the molecular electric field, with special emphasis on the electron-electron component.

Traditional quantum chemistry offers the posibility of solving the electronic wavefunction of the cation $\phi(\mathbf{r})$. From this solution, the cationic electron potential acting on the single-actrive electron V_{e-e} can be computed at all points as

$$V_{e-e}(\mathbf{r}) = \int d\mathbf{r}'\, \frac{|\phi(\mathbf{r}')|^2}{|\mathbf{r} - \mathbf{r}'|}. \tag{5.36}$$

An analytical expression of this two-point integral is not accessible. Quantum chemistry packages allow for an approximated evaluation at points. From these points, we aim to generate models that are capable of accurately fitting the function everywhere. The electrostatic potential (ESP) software was created for this task. We implemented a routine that uses the Psi4 package [100] to obtain the potential at points. Other quantum chemistry models and user-defined values of the potential are also accepted by the ESP package. The calculation of the molecular potential depends on the choice of electronic state. Here, the ground electronic state of the cation is assumed for setting up the calculations. The ESP software aims to create an interpolation of the potential values that allows for a quick evaluation. Afterwards, the cationic electric field can be computed using automatic differentiation. We explore a range of alternatives for function approximation, from quantum-chemistry motivated to pure data-driven approaches. Here, we briefly discuss the alternatives and provide an example of their use[5].

The radial part of the electronic function is often solved in terms of gaussians centered at the nuclei. Introducing the gaussian expansion in Equation (5.36), we obtain integrals that are similar to those generating the error functions. By expanding in terms of M error functions for each of the N nuclei and fitting the coefficients of the expansion, the

---

[5] A potential future addition to the set of models to fit the ESP is the induced polynomial set, as described in Chapter 4.

potential is approximated by

$$\hat{V}_{e-e}(\mathbf{r}) = \sum_{c=1}^{N} \sum_{i=1}^{M} A_{ic} \frac{\mathrm{erf}(B_{ic} d_c)}{d_c}, \tag{5.37}$$

where c labels the nuclei of the molecule, $d_c = |\mathbf{r} - \mathbf{R}_c|$ is the distance to the nucleus c, and A_{ic} and B_{ic} are the parameters to be fitted. In this construction, the number of fitting parameters scales linearly with the number of atoms N.

The asymptotic behaviour of the error functions can be improved by the addition of functions with the correct long-range decay. At long distances, the electronic potential decays as $\sim \frac{1}{r}$. To ease the obtention of this decay, we introduce functions that correctly model it. The potential can be approximated by

$$\hat{V}_{e-e}(\mathbf{r}) = \sum_{c=1}^{N} \left(\sum_{i=1}^{M} A_{ic} \frac{\mathrm{erf}(B_{ic} d_c)}{d_c} + \sum_{j=1}^{K} \frac{Z_{jc}}{d_c + s_{jc}} \right), \tag{5.38}$$

where K is the number of $\frac{1}{r}$ functions, s_{jc} are positive constants that allow for avoiding the divergence at small distances, and Z_{jc} resemble charges. Once more, the number of parameters of this approximation scales linearly with the number of atoms.

Neural networks have been proven to be universal approximators of functions of any dimensionality [167]. The interpolation properties of neural networks have been found to be better comparable to those of other numerical methods. For this reason, the application of neural networks for this approximation problem is well-suited. By having a neural network with activation function σ and N hidden linear layers of weight matrices $\mathbf{A}^{(i)}$ and biases $\mathbf{b}^{(i)}$, where i refers to the layer, the potential can be approximated as

$$\hat{V}_{e-e}(\mathbf{r}) = \mathbf{A}^{(N)} \sigma(\ldots\ (\sigma(\mathbf{A}^{(1)}\mathbf{r} + \mathbf{b}^{(1)}))) + \mathbf{b}^{(N)}. \tag{5.39}$$

The main drawback of this approach is that the asymptotic behaviour of the potential is not guaranteed and, in general, the approximation error of neural networks increases in the extrapolation region with the distance to the closest training point. By construction, the number of parameters in the neural network does not explicitly depend on the complexity of the molecule.

If the set of points used for training is distributed in an equidistant grid, we can use numerical tools such as regular grid interpolators (RGI). These interpolators offer exact evaluations on the training data and a linear fit to the values in between training points. The evaluation time of the RGIs increases with the number of training points, and is formally independent of the number of atoms in the molecule (beside the weak dependence of the grid size on the molecular size). While the approximation of RGIs in dense grids achieves high accuracy, the efficiency of evaluation is highly reduced. This method does not provide extrapolation for performing calculations on points not contained within the hypervolume of the training data. For the evaluation at extrapolation points, which we assume are far enough from the molecule, the potential needs to be defined analytically. We use the expected decay of the complete cationic potential in the long-range regime and define the extrapolation potential as $\frac{1}{r}$. This creates a non-continuous representation of the potential. In practice, this discontinuity does not lead to any major problem.

For any of these choices, the total molecular potential field is given by

$$\hat{V}(\mathbf{r}) = \hat{V}_{e-e}(\mathbf{r}) + \hat{V}_{e-N}(\mathbf{r}), \tag{5.40}$$

where $\hat{V}_{e-N}$ is the electron-nuclei interaction potential. An analytical expression that uses the Coulomb law to model the electron-nuclei interaction can be found in (1.9).

We present two examples of the use of the ESP package: the fitting of the electronic potential of the OCS cation, relevant for LIED calculations, and of the methane cation, a complex molecule in the symmetry group T_d, and of general relevance in chemistry. Both calculations were performed for their equilibrium geometries. The OCS molecule was aligned with the z dimension and its geometry was $\mathbf{x}_{\mathrm{O}} = [0, 0, 3.179]$, $\mathbf{x}_{\mathrm{C}} = [0, 0, 0.989]$ and $\mathbf{x}_{\mathrm{S}} = [0, 0, -1.961]$ in atomic units. The position of the atoms in methane was $\mathbf{x}_{\mathrm{C}} = [0, 0, 0]$, $\mathbf{x}_{\mathrm{H}_1} = [2.05, 0, 0]$, $\mathbf{x}_{\mathrm{H}_2} = [-0.6833, 1.9369, 0]$, $\mathbf{x}_{\mathrm{H}_3} = [-0.6833, -0.9685, 1.677]$ and $\mathbf{x}_{\mathrm{H}_4} = [-0.6833, -0.9685, -1.677]$ in atomic units.

We compute the potential field at points using Psi4 with the unrestricted Hartree-Fock method and the aug-cc-pVTZ basis set. For OCS, the potential training points were computed in an equidistant grid in the space $[-1.5, 1.5] \times [-1.5, 1.5] \times [-4.5, 4.5]$ with a separation of 0.1, 0.1 and 0.05, respectively, giving a total of 162000 points. For methane, the potential was computed in an equidistant grid in the space $[-3, 3] \times [-3, 3] \times [-3, 3]$ with a separation of 0.09 in all dimensions, giving a total of 300763 points. Both grids were set up such that the revelant molecular structure is contained inside of them. The approximations were validated in a grid that extends further from the nuclei and that is denser, i. e., the validation grid samples both the extrapolation and interpolation regimes. For OCS, the validation grid is $[-2, 2] \times [-2, 2] \times [-5, 5]$ with a separation of 0.09, 0.09 and 0.04, respectively, giving a total of 506250 points. For methane, the validation grid was $[-3.2, 3.2] \times [-3.2, 3.2] \times [-3.2, 3.2]$ with a separation of 0.085 in all dimensions, containing 438976 points in total. The RMSEs for each of the fits in the validation grids are shown in Table 7.

Molecule	**erf**	**1/r + erf**	**NN**	**RGI**	**V = 0**
OCS	0.0699	0.0420	0.1844	0.1051	2.4385
Methane	0.0232	0.0221	0.0242	0.0163	0.8119

Table 7: RMSE values for different electrostatic potential fitting methods against Psi4 values (au) on a validation grid of 506250 points for OCS and 438976 points for methane. For referencing the accuracy of these calculations, the RMSE of the naive fit $V = 0$ is added as an extra column.

Comparing to the baseline naive approximation $V = 0$, all the presented methods present orders of magnitude enhanced accuracy. This improvement is significant since the fitting $V = 0$ performs well in the long-range region, which ocuppies a big portion of the grid. The best interpolation and extrapolation properties in both cases are achieved by the error functions or the combination of error functions with $\frac{1}{r}$-like functions. RGI and neural networks provide a fitting that achieves comparable accuracy to the quantum-chemistry-based methods, specially for the interpolation regime. RGI struggle to capture the OCS potential, whose gradient is greater in space, and captures the methane potential with high accuracy. Given that electron trajectories in LIED cover long distances from the molecule, the use of these pure numerical methods for the LIED simulations is not well-suited. We ilustrate the accuracy of the fitted electrostatic potential in Fig. 5.6, in which we present a two-dimensional cut of the error of the fitted potential of OCS

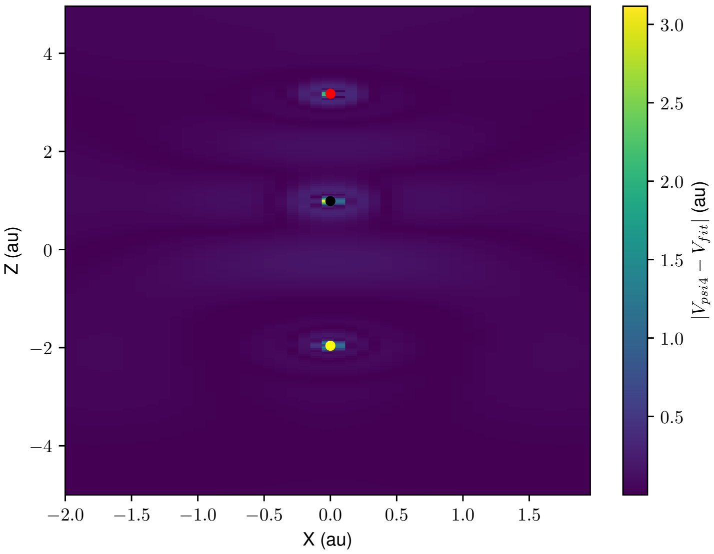

Figure 5.6: Two-dimensional cut of the error in the electrostatic potential. The error is computed as the absolute value of the Psi4 calculated potential minus the fitted potential using error functions and inverse distance functions. The nuclei are contained in the plane and are depicted as red, black and yellow circles that represent oxygen, carbon and sulfur atoms, respectively. The extension of the plane includes extrapolation and interpolation of the training data.

using error functions and $\frac{1}{r}$ functions. The two-dimensional plane of the cut contains the three nuclei that constitute OCS. The fitting error is mainly located very near to the atoms, where the electric field is also larger.

For LIED, the ultimate focus of this calculations is to compute the molecular electric field at all points. The electric field can be computed from the potential field as

$$\hat{\mathbf{E}}(\mathbf{r}) = -\nabla \hat{V}(\mathbf{r}), \tag{5.41}$$

where ∇ refers to the gradient in the Euclidian coordinates. Given that the ESP fitting provides an analytical representation, the electric field can be computed using automatic differentiation techniques. To perform the automatic differentiation, we use the open package JAX[168].

We compare the values of the electric field computed by automatic differentiation of the ESP models with those obtained from Psi4, which internally uses finites diferences. We compute the RMSE for each dimension over the validation points for the electric field of OCS and methane. These errors are tabulated in Table 8. We include an extra column containing the evaluation time of the electric field on the validation points of the different ESP models (excluding the initial compilation time of the automatic differentiation).

We observe that the fittings that use error functions also provide the lowest errors in the approximation of the electric field, consistently with what was observed in the case of the potential field. As the Psi4 electric field is computed using finite differences of the potential field, numerical error might be included in the calculation of validation points. This potentially explains why the electric field of methane for the 1/r + erf fit generates larger error than the erf and neural network, contradicting the potential field results.

Molecule	Method	RMSE [X, Y, Z] (au)	Evaluation time (s)
OCS	erf	[0.5143, 0.5161, 0.5255]	0.83
	1/r + erf	[0.2021, 0.2038, 0.1934]	1.14
	NN	[0.5933, 0.5701, 0.5233]	1.68
	RGI	[0.3346, 0.3328, 0.2737]	1.36
	$\mathbf{E} = 0$	[24.922, 24.919, 17.178]	-
Methane	erf	[0.0561, 0.0559, 0.0556]	1.19
	1/r + erf	[0.1152, 0.1150, 0.1160]	1.19
	NN	[0.0695, 0.0691, 0.0697]	1.59
	RGI	[0.0420, 0.0411, 0.0417]	1.35
	$\mathbf{E} = 0$	[3.7858, 3.7305, 3.6969]	-

Table 8: RMSE in each of the coordinates of the fitted electric field for different methods against the Psi4 values (au) on a validation grid of 506250 points for OCS and 438976 points for methane. The evaluation time (after compilation) of each method for the validation points is included as an extra column. For referencing the accuracy of approximations, we add an extra row with the RMSE of the naive approximation $\mathbf{E} = 0$.

The evaluation time of the nuclear-center-based approximations scales with the number of atoms, whereas the neural network and RGI is independent of the molecular size. For smaller molecules, i. e., less than 10 atoms, the error function approaches also provide the most efficient calculations. This is an important observation when choosing the model for the LIED propagations, where the electric field is evaluated many millions times.

We provide a visualization of the electric fields of the OCS and methane cations. For creating the illustration, the error plus inverse-distance functions fit was used for both molecules. Fig. 5.7 illustrates the electric field in a two-dimensional plane containing some of the nuclei. Additionally, a visualization of the three-dimensional electric field located at random points is depicted. In both molecules, the nuclear electric field is the main component and the field lines roughly point in the radial direction from the atoms. The nodal planes of the electric field of molecules are easily distinguishable in the two-dimensional representations.

The potential field in the ESP framework is fitted using a fixed position of the nuclei. In the LIED calculations, the evaluation of the molecular electric field for rotated nuclear configurations is required. Under a rotation by the Euler angles Ω that generate the rotation matrix $\mathbf{R}(\Omega)$, the electric field at all points can be computed from the initially fitted electric field $\hat{\mathbf{E}}(\mathbf{r})$. By concatenating rotations of the evaluation position and electric field, the electric field is given by

$$\mathbf{E}(\mathbf{r}, \Omega) = \mathbf{R}(\Omega) \cdot \hat{\mathbf{E}}(\mathbf{R}^{-1}(\Omega) \cdot \mathbf{r}). \tag{5.42}$$

5.5.5 Propagation: Runge-Kutta, Kepler orbits and asymptotic behaviour

Initial conditions of electron trajectories can be created using Monte Carlo rejection sampling from the ionisation rate. For a bunch of N photoelectron trajectories, the initial conditions are given by the set

$$\left\{\mathbf{r}_0^i, \mathbf{p}_0^i, \Phi_0^i; t_0^i\right\}_{i=1}^N . \tag{5.43}$$

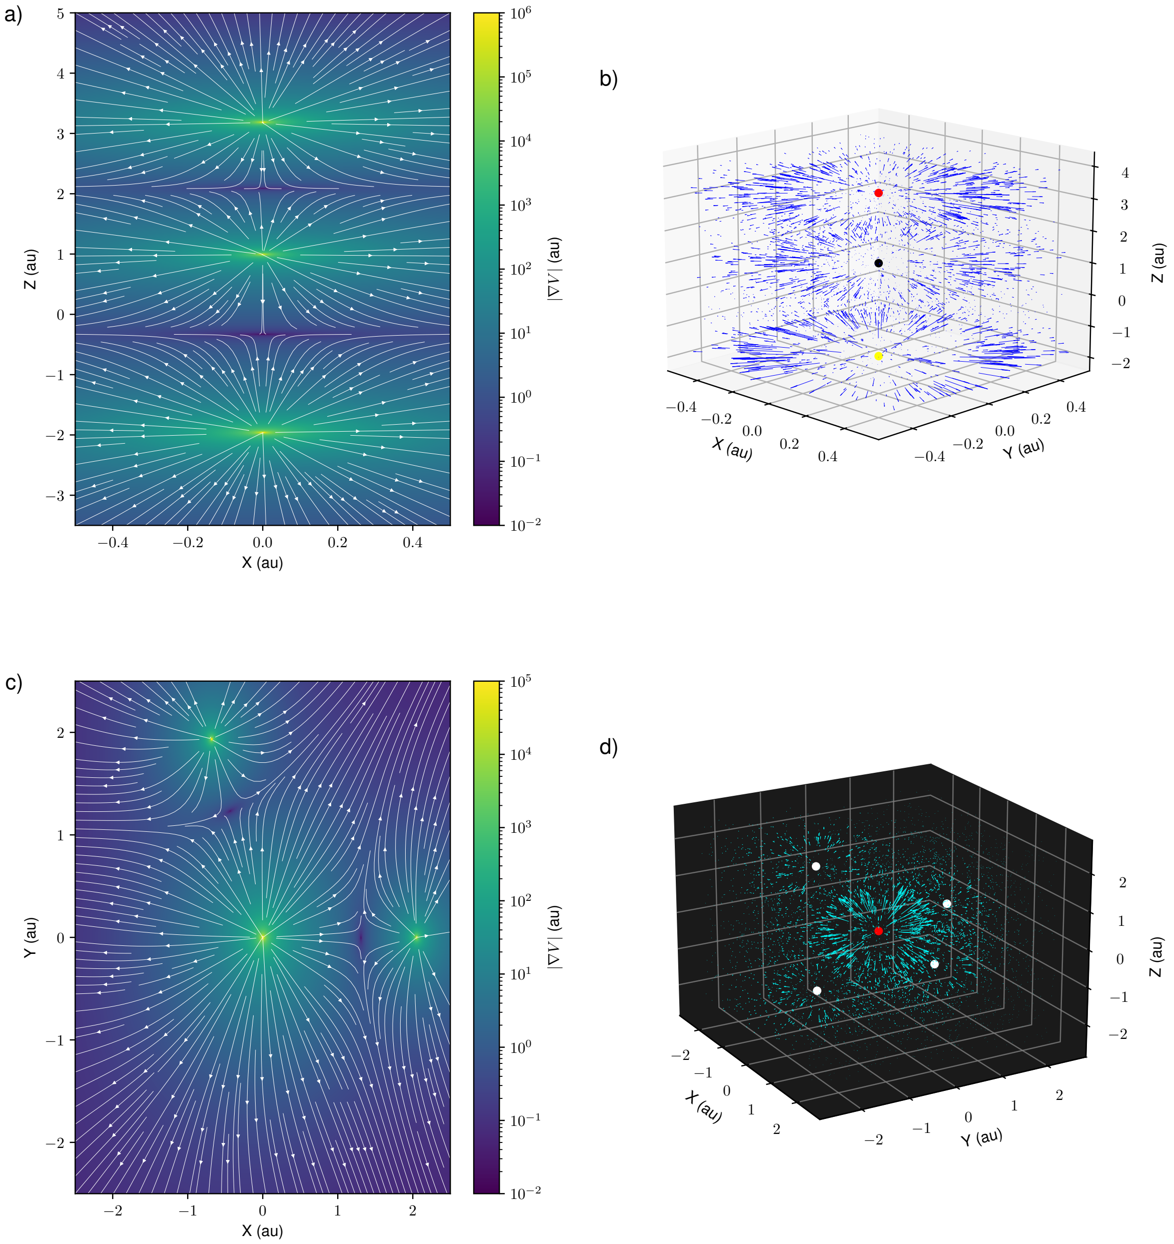

Figure 5.7: Visualization of the fitted electric field over space: **a)** Two-dimensional cut of the electric field in the plane containing the nuclei of OCS. The norm of the electric field at points is depicted by the colorbar. The lines show the direction of the vector field. **b)** Three-dimensional representation of the electric field of OCS on randomly distributed points. The length of each arrow is proportional to the intensity of the electric field. **c)** Two-dimensional cut of the electric field of methane in a plane containing the carbon atom and two of the hydrogens. This plane is equidistant to the two other hydrogens, therefore, their contribution to the electric field in this plane cancels. **d)** Three-dimensional representation of the electric field over space of methane on randomly distributed points. Due to the big difference in the charge of the carbon and the hydrogens, the radial electric field from the central atom is the main contribution to the total electric field.

The laser field in (5.21) finishes its action at time $t = T/2$. The propagation step under the action of the laser field aims to compute the quantities at this time, $\left\{\mathbf{r}^i(T), \mathbf{p}^i(T), \Phi^i(T)\right\}_{i=0}^N$. The asymptotic properties of photelectrons can be calculated using these values. In other words, the propagation aims to find the solution of the system of coupled differential equations

$$\begin{aligned}
\mathbf{r}^i(T) &= \mathbf{r}_0^i + \int_{t_0}^{T} dt\ \mathbf{p}^i(t) \\
\mathbf{p}^i(T) &= \mathbf{p}_0^i - \int_{t_0}^{T} dt\ \mathbf{E}(t) \\
\Phi^i(T) &= \Phi_0^i - \int_{t_0}^{T} dt\ \left(\frac{|\mathbf{p}^i(t)|^2}{2} - \frac{2\,Z}{|\mathbf{r}^i(t)|}\right),
\end{aligned} \tag{5.44}$$

where $\mathbf{E}$ is the total electric field, and the differential equation for the classical phase was previously defined in Equation (5.13).

This set of equations defines a system of coupled partial differential equations. We resort to numerical methods to find approximations of their solution. In particular, we choose Runge-Kutta methods— both implicit and explicit propagators— and use their implementation in the Jax-based Diffrax library [169]. The implict Runge-Kutta method is the main focus of chapter 6 and a description of the method is given in section 6.5.

Trajectories that scatter close to the nuclei are of utter importance to reconstruct the molecular geometry. Traditional propagators (Runge-Kutta) find a loss of efficiency for such trajectories as the rates of change in this regime are large. For this reason, decreasing time steps sizes are required for obtaining acceptable accuracies. The computational time of the propagation of one time step is roughly independent on its step size. Therefore, the reduction in time step size hinders the computational efficiency of the semiclassical model.

To avoid trajectories getting computationally "stuck" when approaching the nuclei, we introduce the Kepler orbit approximation. When an electron approaches the i-th nucleus of the molecule, the nuclear electric field becomes orders of magnitude larger than the laser field. Then, at some distance r^i_{eff}, the magnitude difference is so large that the laser field can be perceived as a perturbation to the nuclear field. The potential of the central force for the i-th nuclei can be approximated by

$$V^i(r) = -\frac{Z^i_{\text{eff}}}{r}, \tag{5.45}$$

where Z^i_{eff} is the effective charge of the nuclei which arises from the screening of the core electrons to the nuclear charge Z^i, satisfying $Z^i_{\text{eff}} < Z^i$. The effective charge depends on the distance to the nuclei, as in closer distances less core electrons are percieved. In practice, a fixed effective charge was chosen by calculating the mean effective charge over distances up to r^i_{eff}.

In our implementation, the user selects the ratio between the nuclear field and the peak laser field for which the laser can be considered as a perturbation (this ratio defaults to 1000). From this, the effective radius for each atom in the molecule is computed and is used to calculate the effective charge.

Neglecting the contribution of the laser field, the dynamics of electrons in this regime are equivalent to orbital mechanics, which offer an analytical solution. This analytical solution is used to find the outgoing momentum and position of the electron after scattering. The laser field perturbation is added *a posteriori* to the outgoing momentum. The additional momentum $\Delta\mathbf{p}$ is given by

$$\Delta\mathbf{p} = -\int_{t_{\text{in}}}^{t_{\text{in}}+\Delta t} dt\ \mathbf{E}_{\text{laser}}(t) = \mathbf{A}_{\text{laser}}(t_{\text{in}} + \Delta t) - \mathbf{A}_{\text{laser}}(t_{\text{in}}), \tag{5.46}$$

where Δt is the propagation time of the scattering event and t_{in} is the time where the Kepler orbit started.

The complete stategy for the propagation is as follows: propagate every trajectory until either the finalization of the laser field or finding the electron at a distance to a nuclei smaller than the respective effective radius. Store the finished trajectories. Compute the Kepler orbit approximation on the rest. Repeat the first step until all trajectories are finished or until a maximum number of Kepler iterations is reached.

When the laser field is finished, the force becomes time-independent and, from this point, energy is conserved. Electrons with a kinetic energy that exceeds the potential energy barrier will become ionised, while those who don't, recombine with the molecule. Kepler orbital mechanics equations can be once more applied to calculate the asymptotic motion of ionised electrons. The asymptotic momentum $\mathbf{p}_\infty$ is the quantity of interest, as it is closely related to the observable of the experiment. The asymptotic momentum can be computed using the momentum and position of photoelectrons at the end of the laser.

5.5.6 Detection

Electrons follow a free-particle trajectory as they escape the molecular field. After a time much longer than the characterisitic electronic timescales, photoelectrons are detected using a velocity map imaging (VMI) spectrometer [170, 171]. At a basic level, a VMI consists of two parallel plates with different voltages that generate a uniform electric field perpendicular to the plates, and two detectors in both directions of the field, one for detecting anions and one for cations. Escaping electrons are accelerated as a result of the uniform field, leading to the loss of the asymptotic momentum information along this axis. The VMI electric field direction is oriented perpendicular to the laser field axis, as the momentum component along the laser axis is the key observable for the imaging technique.

The detection of photoelectron velocity in the VMI goes as follows. All photoelectrons are forced to collide with the detector by the action of the uniform field. The source from where photoelectrons originate is aproximately a point source in the distance scales of the detector. Therefore, the measured perpendicular distance to the source is proportional to the perpendicular momentum. The perpendicular momentum of the electrons can be estimated by using the plates-to-detector distance and the generated electric field strength. The generated observable from the VMI is the PAD, that measures the probability density of finding an electron with respect to the perpendicular asymptotic momentum $\mathbf{p}_\infty^\perp$.

Due to the discrete nature of the representation of variables in the semiclasical model, we calculate the theoretical PAD as a histogram over bins of electrons with a similar asymptotic perpendicular momentum. In practice, this histogram closely resembles the pixel resolution of the detector. If the width of the bins equals the resolution of the detector, no error is introduced at this step.

The histogram is not directly created as the sum of electrons. The asymptotic phase that each electron carries Φ^j_∞ needs to be taken into account. Two electrons that end up in the same bin can produce a constructive or a destructive sum, depending on their relative phases. For a LIED simulation of N electrons, the bin centered at momentum $\mathbf{p}$ with width $\Delta\mathbf{p}$ has an intensity given by [152]

$$I(\mathbf{p}) = \left| \sum_{j=0}^{N} \exp(i\ \Phi^j_\infty)\ \Theta\left(\left| [\mathbf{p}^\perp_\infty]^j - \mathbf{p} \right| - \Delta\mathbf{p} \right) \right|^2 , \tag{5.47}$$

where $[\mathbf{p}^\perp_\infty]^j$ is the perpendicular momentum of the j-th electron and Θ is the Heaviside step function, such that if the asymptotic perpendicular momentum is not contained in the box of size $\Delta\mathbf{p}$, the electron is not considered in the calculation.

6 Learning phase-space flows using time-discrete implicit Runge-Kutta PINNs [6]

In this chapter, we use discrete-time implicit Runge-Kuta (IRK) PINNs of very high-order to create an algorithm that is capable of simultaneously learning all possible trajectories of systems that have explicitly time-independent and/or periodic forces.

This project originated as a possibility to construct an efficient and accurate propagator of the trajectories of photoelectrons in LIED, in contrast to applying standard Runge-Kutta methods iteratively at each time step. However, our approximation of the force field in LIED is not time-independent nor periodic, given that the envelope of the laser field breaks its periodicity. For this reason, the application of this method in LIED was limited to very specific experimental setups.

While limited in its application to LIED, this work is interesting on its own. It provides a summary of the trending PINNs approach, and of its application in the, often overlooked, time-discrete version of PINNs to different systems. Additionally, this application allows for the creation of numerical propagators that are capable of solving all trajectories in the domain at the same time.

6.1 Introduction

Physics-Informed Neural Networks (PINNs) have emerged as a prominent and dynamic area of research for solving differential equations [18, 172], for example, for modeling the physics of fluid dynamics [173–175] or quantum mechanics [176]. Unlike traditional neural networks, that learn solely from data, PINNs use both data and physical equations to guide the learning process [177–179].

PINNs can be effectively employed using continuous and discrete representations of time. The time-continuous approach uses space and time variables as inputs, and learn to satisfy the differential equations across the entire domain of interest. This can be impractical without data distributed across multiple time slices. In addition, time-continuous PINNs also encounter difficulties with high-frequency oscillations and stiff problems, lacking a clear strategy to deal with them. On the other hand, the discrete-time PINNs learn to model changes within a fixed discrete time step, utilizing only spatial information from a single time slice. This approach improves the accuracy in solving stiff problems [180] by leveraging the A-stability of implicit Runge-Kutta (IRK) methods [181]. This is especially significant for tackling the stiff problems prevalent in particle trajectory simulations and Differential Algebraic Equations (DAEs), known for their inherent infinite stiffness [181, 182].

In this paper, we extend the discrete-time IRK-PINN scheme [18] by generalizing it to a larger number of dimensions in both input variables and output quantities. Importantly, we have adapted this approach to develop a new efficient numerical scheme tailored to finding phase-space flows of classical particle trajectories, described by second-order differential equations of motions. The new algorithm enables the simultaneous learning

---

[6] This chapter is based on the publication by Álvaro Fernández Corral, Nicolás Mendoza, Armin Iske, Andrey Yachmenev, Jochen Küpper, ”Learning phase-space flows using time-discrete implicit Runge-Kutta PINNs”, [arXiv:2409.16826] [cs.LG] (2024), published in the International Conference on Scientific Computing and Machine learning SCML2024. My contributions to this work were the conceptualization of the work, creation and preparation of code in an open repository, performance of calculations, generation of figures and preparation of the manuscript.

of all possible trajectories within the specified phase-space and is particularly suited for explicitly time-independent and time-periodic forces.

The manuscript is organized as follows: The details and application of the IRK-PINN scheme to phase-space flows are described in subsection 6.2. Subsection 6.3 presents some illustrative application examples with time-independent and periodic forces. The review and application details of the IRK-PINN scheme, including its successful application to a number of various differential equations, are elaborated in subsections 6.5.1–6.5.3.

6.2 Phase-space flows with IRK-PINNs

In applications to physics problems, the IRK-PINN method was primarily applied to functions $\mathbf{u} : \mathbb{R}^{1+d} \rightarrow \mathbb{R}^m$ which represent background "fields" of the system, such as the temperature, pressure, velocity flow, heat convection, etc., see subsection 6.5.1 for details. The literature on applying PINN-based IRK schemes to particle trajectory analysis in physical systems is limited, with the exception of work focusing on first-order differential equations [183]. We are interested in solving differential equations of the general form

$$\mathbf{F}(\mathbf{x}, \dot{\mathbf{x}}, \ddot{\mathbf{x}}, \dots, t) = \mathbf{0}, \tag{6.1}$$

which encompass, for example, Newton's equation, expressed as $\ddot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \dot{\mathbf{x}}, t)$. Our goal is to adapt and apply the IRK-PINNs scheme to effectively compute particle flows $\mathbf{x}(t)$ resulting from a force $\mathbf{f}(\mathbf{x}, \dot{\mathbf{x}}, t)$. We begin by transforming the second-order Newton-type differential equation into a system of first-order differential equations using the $2d$-dimensional phase-space coordinates, denoted as $\chi = (\mathbf{x}, \dot{\mathbf{x}})$, with the dimensionality d of the vector $\mathbf{x}$. We define our function of interest as $\mathbf{u}(t, \chi_i) = \hat{\chi}_i(t)$, with $\hat{\chi}_i : \mathbb{R} \rightarrow \mathbb{R}^{2d}$ representing the flow function. This flow function defines the trajectory of any point in phase-space – a one parameter curve in our phase-space manifold – that satisfies the condition $\hat{\chi}_i(t_n) = \chi_i$. By adhering to the equations of motions, this condition ensures a unique solution. Indeed, since the value of $\mathbf{u}(t_n, \chi_i)$ at the initial time t_n is $\mathbf{u}_n(\chi_i) = \chi_i$, we can determine the phase-space values at time t_{n+1} by employing the IRK-PINNs time propagation method. By applying the IRK scheme, see (6.5) in subsection 6.5.1, using a fully connected feedforward neural network (FNN) also known as a multilayer perceptron (MLP), to the general trajectory equation (6.1), we thus determine the trajectories that satisfy the following set of differential equations

$$0 = \mathcal{N}[\chi] + \frac{d}{dt}\chi \equiv \begin{pmatrix} \dot{\mathbf{x}} \\ \mathbf{f}(\mathbf{x}, \dot{\mathbf{x}}, t) \end{pmatrix} + \frac{d}{dt}\begin{pmatrix} \mathbf{x} \\ \dot{\mathbf{x}} \end{pmatrix}. \tag{6.2}$$

Any additional boundary conditions that may be essential for the well-defined nature of the problem can also be accommodated.

The algorithm's strength lies in its ability to predict the future state of every point in the phase-space at time t_{n+1} by learning from only a limited sample of phase-space data at the initial time t_n. In addition, as a high-order IRK algorithm, it is well suited to solve stiff problems, enhancing computational efficiency in the propagation of numerous trajectories.

Moreover, this algorithm is particularly effective for forces that are explicitly time-independent and/or periodic in time. If the force doesn't explicitly change over time and only relies on the position or velocity of the particle, both being implicitly time-dependent vector fields, the neural network can effectively predict the trajectories by learning from a training set in the phase-space. For periodic forces, we impose the time-propagation step size to match the period of the force $\Delta t := T$, so that the neural network output $\chi(t_{n+1})$ would reside on a phase-space manifold identical to the initial one at t_n. The periodic nature of the force is essential, as the phase-space manifold is entirely dependent on the force governing the particles' motions, necessitating that $\mathbf{f}(\mathbf{x}, t_n + \Delta t)$ remains equivalent to $\mathbf{f}(\mathbf{x}, t_n)$. In both cases, or a combination, the phase-space manifold remains unchanged after a single time step propagation. Consequently, to predict trajectories for subsequent time slices, the IRK-PINN algorithm can be recurrently applied after training.

6.3 Results

6.3.1 Keplerian orbits

A massive body under the action of a central potential moves according to Newton's equation of gravity. The trajectories that orbiting objects follow are called Keplerian orbits. The analytical solution of N-body problems is complicated due to the implicit time-dependency of the force $F \sim 1/r_{ij}^2(t)$, arising from the varying distances between different objects $r_{ij}^2(t)$. For this reason, low-order IRK schemes with small time stepping were used to propagate these systems [184]. Following the discussion in the previous section, our phase-space flows approach offers an efficient alternative for propagating these explicitly time-independent systems. This is achieved by recurrently applying the neural network, requiring only a single training session of the IRK-PINN and thus eliminating the need for additional approximations.

We focus on central forces, which conserve angular momentum, leading to the confinement of trajectories within a two-dimensional plane defined by the initial momentum and the radial vector [185]. Initially, our efforts were directed at solving the fixed-central-mass problem in a two-dimensional Euclidian space with a Coulomb potential. However, after experimenting with various MLPs and parameters, we observed that the algorithm was unable to accurately predict the trajectories for this problem. We believe that this limitation arises from the fact that not all points in the initial phase-space correspond to physically viable solutions. Given that the algorithm attempts to simultaneously solve for all trajectories, these divergent paths likely cause a general diverging behaviour of the method.

The accurate and efficient calculation of Keplerian orbits is critical in simulations of particle diffraction from the Coulomb potential with a small impact parameter. This is especially important for simulations of laser-induced electron diffraction [134], where the diffraction of highest-energy electrons is influenced by these orbital dynamics [138, 186, 187].

To avoid singularities, we replaced the Coulomb potential with a central potential characterized by a Gaussian function $V(\mathbf{x}) = -V_0 \exp(-\mathbf{x}^2/a^2)$, where V_0 is the depth of the potential well and the a parameter determines its extent. This choice of potential is a standard first-order approximation to the Coulomb potential in strong-field physics [188] and provides Kepler-like orbits, while not provoking diverging trajectories, due to its finite value at the origin. Defining the force as the negative gradient of this potential,

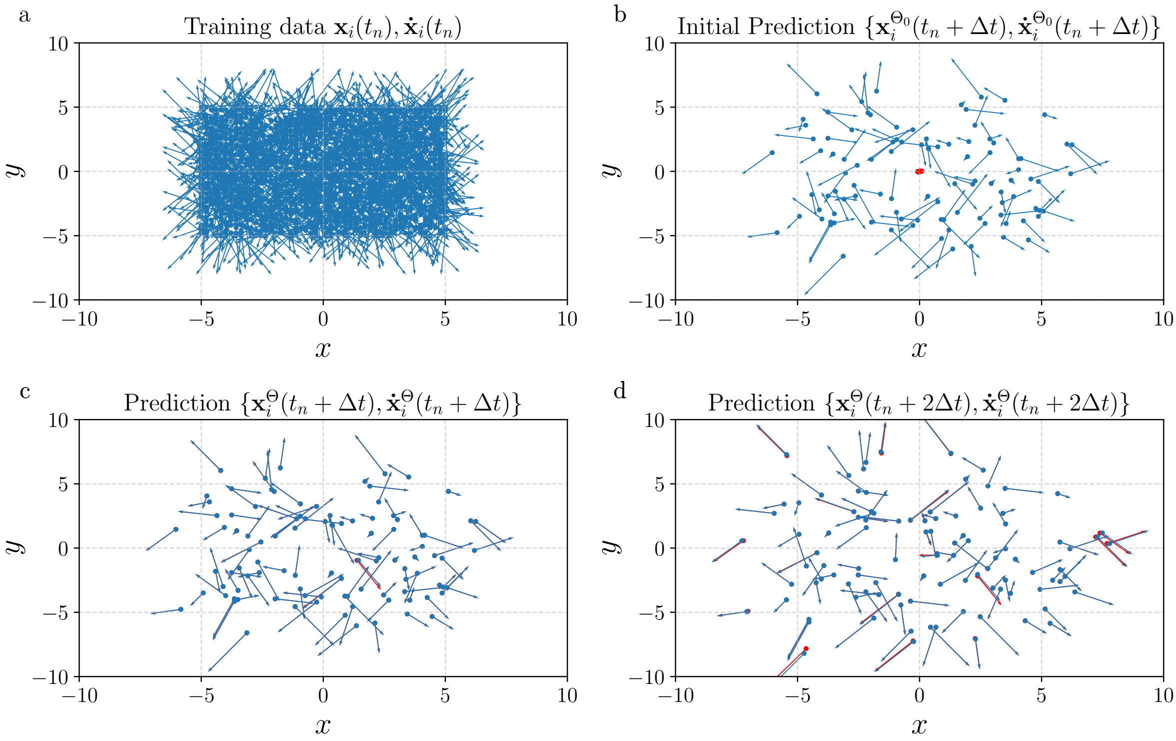

Figure 6.1: Calculated trajectories of a particle in a central Gaussian potential, with potential depth $V_0 = 10$ and extension $a = \sqrt{5}$. The IRK-PINN and numerically exact solutions are plotted with blue and red colours, respectively. Panel (a) shows the entire training phase space consisting of 2 000 points. Due to the complex distribution of the solutions with respect to the initial position, a big set of training data is needed. Panel (b) shows the result of the low-order IRK method (blue) on the validation set, consisting of 100 points, for visual clarity. In addition, the PINN guess on parameter initialisation is displayed (red). Panels (c) and (d) show the IRK-PINN predictions on the validation set for the time step Δt and $2\Delta t$, respectively. The IRK-PINNs approach adopted a Runge-Kutta order $q = 100$, time step $\Delta t = 0.8$, dense MLP containing 5 hidden layers with 200 nodes each, and a bipolar sigmoid as the activation function. The training was carried out for 100 000 epochs with the ADAM optimiser followed by 40 000 epochs with L-BFGS-B. Total L^1-error = 0.58 %.

the equations of motion become

$$\mathcal{N}[\chi] = -\begin{pmatrix} \dot{\mathbf{x}} \\ -\nabla_{\mathbf{x}} V(\mathbf{x}) \end{pmatrix} \equiv -\begin{pmatrix} \dot{\mathbf{x}} \\ -2\frac{V_0}{a^2}\exp(-\frac{\mathbf{x}^2}{a^2})\mathbf{x} \end{pmatrix}. \tag{6.3}$$

The domain of the model is $(\mathbf{x}, \dot{\mathbf{x}}) \in \Omega = \mathbb{R}^2 \times \mathbb{R}^2$. Considering that the phase-space flows approach uses an equal number of phase-space coordinates and objective functions, we will work within a 4-dimensional phase-space and output space. For benchmarking the IRK-PINN results, we generated accurate solutions using low-order adaptive IRK schemes, specifically Kværnø's 5/4 method [158], as implemented in the Diffrax library [169].

The accuracy of the IRK-PINNs solution, along with the details of the neural network, are demonstrated in Fig. 6.1. It is important to note that the central force in this model is explicitly time-independent. This aspect, combined with comprehensive solution for the entire phase space, enables us to propagate in time steps of Δt by recurrently applying the model at each time step. To illustrate this, in Fig. 6.1 d we show the accuracy of predictions of the neural network at a doubled time step $2\Delta t$. In addition, this panel shows the ability of the PINN to predict trajectories from points that are not contained inside the training region.

6.3.2 Charged particle under the action of a sinusoidal laser

In this example, we consider the motion of a charged particle under the influence of a periodic external electric field, i. e., a laser field. The force exerted on a particle with charge q by an external field $\mathbf{E}(\mathbf{x}, t)$ is given by $\mathbf{F}(\mathbf{x}, t) = q\mathbf{E}(\mathbf{x}, t)$. As in the previous example, we continue to work within a 4-dimensional phase space defined by $(\mathbf{x}, \dot{\mathbf{x}}) \in \Omega = \mathbb{R}^2 \times \mathbb{R}^2$, also corresponding to a 4-dimensional output space.

For simplicity, we assume that the wavelength of the laser field is much larger than the scale of particle movement. The electric field, characterised by an angular frequency ω and incident at an angle α relative to the x-axis, is represented by the function $\mathbf{E}(\mathbf{x}, t) = (E_x(t), E_y(t)) = (A\cos(\omega t)\cos(\alpha), A\cos(\omega t)\sin(\alpha))$, where A denotes field's amplitude. Choosing the unit system $M = q = 1$, the differential equation of motion can be expressed as

$$\mathcal{N}[\chi] = -\begin{pmatrix}\dot{\mathbf{x}}\\ \mathbf{E}(\mathbf{x})\end{pmatrix} = -\begin{pmatrix}\dot{x}\\ \dot{y}\\ A\cos(\omega t)\cos(\alpha)\\ A\cos(\omega t)\sin(\alpha)\end{pmatrix}. \tag{6.4}$$

This particular form of the laser field is chosen for its analytical solvability, substantive complexity, and periodic behaviour in time. Given any initial conditions $(x_0, y_0, \dot{x}_0, \dot{y}_0)$ at time t_0, the analytical solution can be readily obtained. We examine the system with a laser of period $T = 1 := \Delta t$, amplitude $A = 10$ and incident angle $\alpha = 0.5$. To benchmark our IRK-PINNs implementation, we compare its results for different Runge-Kutta orders with the analytical solution. The results of the simulations and the details of the neural networks used are presented in Fig. 6.2. These results were obtained for multiple periods of the electric field by recurrently applying the neural network, which was trained only once, for the first period, using a set of training points distributed across a subset of the phase-space. For the purpose of illustration and avoiding overfitting, a validation set contained in a different subset of the phase-space was used to generate the plots.

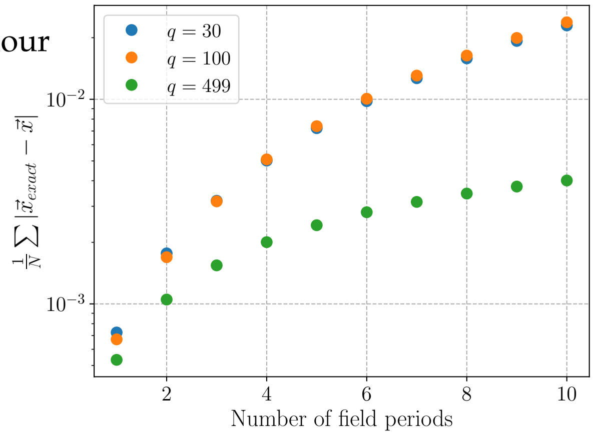

Figure 6.2: Accuracy of IRK-PINNs method for different Runge-Kutta orders $q = 30, 100, 499$ in solving the trajectories of a charged particle in periodic electric field, as described by (6.4) with the parameters $\alpha = 0.5$ and $A = 10$. Accuracy is quantified using the L^1-error calculated on a validation set of 2 500 points. Training involved 50 000 epochs each with the ADAM and L-FBGS-B optimizers. The model architecture included an MLP with 5 hidden layers, each composing 32 nodes, and employed the $SiLU(x)$ activation function. The training utilized 5 000 points.

6.4 Conclusions

We introduced a versatile algorithm designed to effectively solve a broad range of differential equations. The algorithm was validated by generating accurate results for both functional PDEs and equations of motion. Notably, the application of PINNs as a propagator for explicitly time-independent and periodic forces represents a significant advancement over conventional low-order IRK methods.

Further work should focus on addressing the problem of divergent trajectories, particularly in cases like Keplerian orbits under a Coulomb $\sim 1/|\mathbf{x}|$ potential. Overcoming this divergences would particularly enhance the algorithm's utility in solving stiff dynamical systems, such as the N-body problem [184] and the dynamics of charged particles in time-independent or periodic external fields [138, 186].

6.5 Appendix

6.5.1 High-order implicit Runge-Kutta scheme

We start by examining a vector-valued function $\mathbf{u} : \Sigma \rightarrow \mathbb{R}^m$, which is defined over spacetime vectors $(t, \mathbf{x}) \in \Sigma := \mathbb{R} \times \Omega \subset \mathbb{R}^{1+d}$ in the domain formed by a combination of 1-dimensional time and d-dimensional space. This function is defined to be the solution of a set of non-linear coupled differential equations, represented as $\partial_t \mathbf{u} + \mathcal{N}[\mathbf{u}] = 0$, $\forall (t, \mathbf{x})$ in Ω. Additionally, we considered the possibility of incorporating boundary conditions, denoted as $\mathcal{B}_\alpha[\mathbf{u}(\partial\Omega)] = 0$, which could depend on the derivatives of $\mathbf{u}$ at the boundary.

We sought to develop an algorithm that utilizes the dataset $\{\mathbf{x}_k, \mathbf{u}(t_n, \mathbf{x}_k)\}_{k=1}^{N}$ at a specific time slice t_n and the differential equations representing the time evolution to accurately predict $\mathbf{u}(t_{n+1}, \mathbf{x})$ at the next time slice t_{n+1}. To describe the time propagation from t_n to t_{n+1}, we introduced an IRK scheme of q-th order, defined by a set of coupled equations

$$\mathbf{u}_{n+c_i} = \mathbf{u}_n - \Delta t \sum_{j=1}^{q} a_{ij} \mathcal{N}[\mathbf{u}_{n+c_j}] \quad \forall i \in \{1, \ldots, q\}, \tag{6.5a}$$

$$\mathbf{u}_{n+1} = \mathbf{u}_n - \Delta t \sum_{j=1}^{q} b_j \mathcal{N}[\mathbf{u}_{n+c_j}], \tag{6.5b}$$

where $\mathbf{u}_{n+c_i}(\mathbf{x}) := \mathbf{u}(t_n + c_i \Delta t, \mathbf{x})$, $\Delta t = t_{n+1} - t_n$ and a_{ij}, b_j, and c_i are the Butcher-tableau coefficients for a chosen IRK order q. The computation of the Butcher tableau involves expanding the solution into a Taylor series and matching the coefficients to the actual solution up to the desired order of accuracy. The choice of this coefficients is not unique, and it specifies the particular IRK method. We chose the Gauss-Legendre Runge-Kutta method, which is A-stable for all orders [189], to find the coefficients throughout our implementation. However, other methods, such as the Lobatto [190], the Radau [191], or the diagonally implicit Runge-Kutta methods [192], result in different accuracy, stability, and efficiency properties. Testing the pros and cons of using different Butcher-tableau coefficients in applications with IRK-PINNs is beyond the scope of this manuscript, although it should be further explored in the future.

The theoretical analysis of the IRK algorithm suggests that the deviation from the exact result scales as $\mathcal{O}(\Delta t^{2q})$ [189]. Although it may initially appear that the error would grow with increasing q for time steps $\Delta t > 1$, it is important to recognize that this error is dimensionful and requires a constant to render it dimensionless for proper interpretation.

For a chosen IRK order q, we place a neural network prior on all intermediate calculations and the final output

$$\mathcal{U} := (\mathbf{u}_{n+c_1}, \dots, \mathbf{u}_{n+c_q}, \mathbf{u}_{n+1}). \tag{6.6}$$

We observe that $\mathcal{U}(\mathbf{x}) \in \mathbb{R}^{m\times(q+1)}$ can be conceptualized as an $m \times (q+1)$ matrix, where m is the output dimension of $\mathbf{u}$ which is not necessarily equal 1. Consequently, we employ an MLP denoted by $\mathcal{U}^\theta : \mathbb{R}^d \to \mathbb{R}^{m\times(q+1)}$, to closely approximate the desired function $\mathcal{U}$. We define $\mathbf{u}_j^\theta$ in a manner analogous to (6.6) as the parameter-dependent approximation. Subsequently, we define a set of parameter-dependent quantities to be used in the loss function as

$$\mathbf{k}_i^\theta = \mathbf{u}_{n+c_i}^\theta + \Delta t \sum_{j=1}^{q} a_{ij} \mathcal{N}[\mathbf{u}_{n+c_j}^\theta] \quad \forall i \in \{1, \dots, q\} \tag{6.7a}$$

$$\mathbf{k}_{q+1}^\theta = \mathbf{u}_{n+1}^\theta + \Delta t \sum_{j=1}^{q} b_j \mathcal{N}[\mathbf{u}_{n+c_j}^\theta]. \tag{6.7b}$$

Comparing (6.7) with (6.5), it is evident that if $\mathcal{U}^\theta$ effectively approximates $\mathcal{U}$, then the earlier definition implies $\mathbf{k}_i^\theta \approx \mathbf{u}_n$ $\forall i \in \{1, \dots, q+1\}$. Using this relationship, we can express the loss function as

$$\mathcal{L}_\mathcal{N}(\theta, \{\mathbf{x}_k\}_{k=1}^N) = \frac{1}{N} \sum_{k=1}^{N} \sum_{i=1}^{q+1} \left\| \mathbf{k}_i^\theta(\mathbf{x}_k) - \mathbf{u}_n(\mathbf{x}_k) \right\|^2, \tag{6.8}$$

where $\{\mathbf{x}_k\}_{k=1}^N$ is a set of spatial points randomly distributed in the phase-space Ω.

Furthermore, if the differential equations include A boundary conditions of the type $\mathcal{B}_\alpha[\mathbf{u}(\partial\Omega)] = 0$, an additional term can be added to the loss function

$$\mathcal{L}_\mathcal{B}(\theta, \{\breve{\mathbf{x}}_b\}_{b=1}^B) = \frac{1}{B}\frac{1}{A} \sum_{b=1}^{B} \sum_{\alpha=1}^{A} \left\| \mathcal{B}_\alpha[\mathbf{u}(\breve{\mathbf{x}}_b)] \right\|^2, \tag{6.9}$$

where $\{\breve{\mathbf{x}}_b\}_{b=1}^B$ is a set of spatial points distributed in the boundary $\partial\Omega$.

The total loss function used in the optimization is expressed as a weighted sum of the aforementioned terms, as

$$\mathcal{L}(\theta, \{\mathbf{x}_k\}_{k=1}^N, \{\breve{\mathbf{x}}_b\}_{b=1}^B) = \omega_\mathcal{N} \mathcal{L}_\mathcal{N} + \omega_\mathcal{B} \mathcal{L}_\mathcal{B}. \tag{6.10}$$

To ensure best convergence, the ratio of the weights $\omega_\mathcal{N}$ and $\omega_\mathcal{B}$ can be chosen different depending on the problem.

6.5.2 Methodology

We implemented the IRK-PINNs using the JAX [193] and Flax [194] Python libraries and optimized the weights of the PINN by initially approaching the minimum using

Optax [115] with the Adam first-order optimizer [195] and then refined the result by switching to the second-order optimization L-BFGS-B [196] available in the Jaxopt library [197]. For the PINN, we employed a fully-connected dense MLP with different number and structure of hidden layers and types of activation functions, depending on the problem addressed. We found that $SiLU(x) = \frac{x\exp(x)}{1+\exp(x)}$ and bipolar sigmoid $f(x) = \frac{\exp(x)-1}{\exp(x)+1}$ [198] activation functions work best for the problems considered in this study.

For benchmarking, we obtained accurate results using the low-order Runge-Kutta approach from the Diffrax library [169], which provides various implicit and explicit Runge-Kutta methods of different orders.

Our developed model is highly versatile and easy to use. Using just one class type, it successfully handled a range of applications, some of which are highlighted in subsection 6.5.3.

6.5.3 Additional results

In this section, we present the performance of the IRK-PINNs scheme applied to systems that handle fields $\mathbf{u} : \mathbb{R}^{1+d} \to \mathbb{R}^m$, similar to those investigated previously [18]. These differ conceptually from a particle's equations of motion, which deal with coordinates $(\mathbf{x}, \dot{\mathbf{x}}) : \mathbb{R} \to \mathbb{R}^{2d}$, in the manner that fields do not have to follow a direct relation with the phase-space of the system.

In the following, we selected a number of functional PDEs with different input/output dimensions to illustrate the versatility of the algorithm.

Heat Equation in a 2D Plate

We deal with a system with a 2-dimensional input: the heat equation over a surface. This equation describes the evolution of a scalar field $T : \mathbb{R}^{2+1} \to \mathbb{R}$, representing the temperature of a 2D system in time. Using our formalism, the heat equation is given by

$$\mathcal{N}[T] = -c^2 \nabla^2 T, \tag{6.11}$$

where c is the thermal diffusivity, a constant that measures the rate of heat transfer inside the material.

We will deal with the system shown in [199], which is a two dimensional sheet $(x, y) \in \Omega = [0, 2]^2$ with $c = 1/3$ and periodic boundary conditions for the temperature given by $T(t, 0, y) = T(t, 2, y) = T(t, x, 0) = T(t, x, 2) = 0$. We will also impose the initial condition of a heated lower half plane $T(0, x, y) = 50(1 - \Theta(y - 1))$, with Θ being the unit (Heaviside) step function. In essence, this represents a 2×2 square that is initially heated and in contact with a cold resevoir at its boundary with $T = 0$. The solution to this equation is given by

$$\begin{aligned} T(t, x, y) = \sum_{m,n=1}^{\infty} \frac{200}{\pi^2} \frac{(1 + (-1)^{m+1})(1 - \cos\frac{n\pi}{2})}{mn} \cdot \\ \cdot \sin\frac{m\pi x}{2} \sin\frac{n\pi y}{2} e^{-\frac{c^2\pi^2}{4}(m^2+n^2)t}. \end{aligned} \tag{6.12}$$

The results of our IRK-PINN prediction, compared to (6.12) and the description of the used neural network, are shown in Fig. 6.3. The IRK-PINN is able to learn the solution of the heat equation, even for relatively small neural network and sample size.

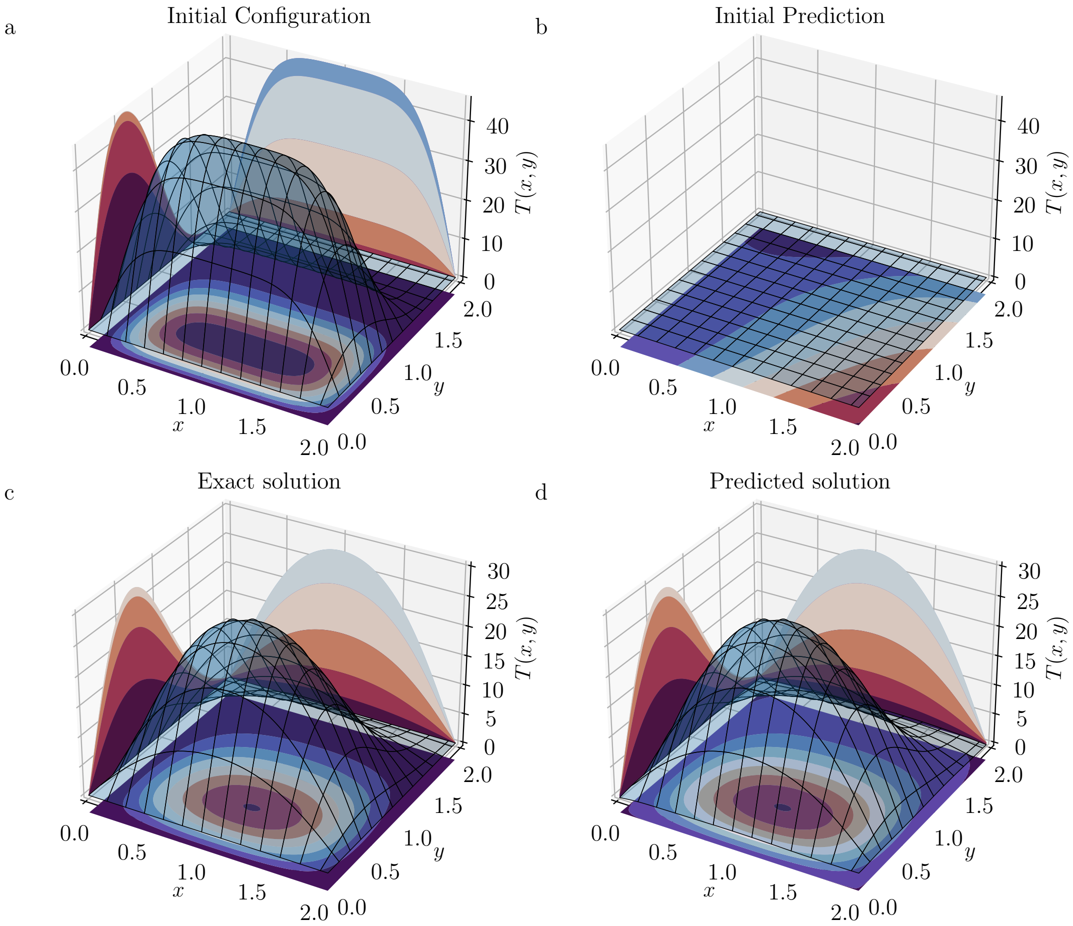

Figure 6.3: Fit to the heat equation data with IRK-PINNs of order $q = 100$ and time step $\Delta t = 0.7$, compared to the analytical solution of the problem. The training was performed for 25000 epochs with ADAM, 40000 epochs with L-BFGS-B, using an MLP with 20 hidden layers with 32 nodes each and $SiLU$ as activation function. (a) The initial configuration space, consisting of $N = 500$ points. (b) The predition of the PINN on parameter inisialization. This initial prediction is totally arbitrary and far from the correct result, which shows the robustness of the learning process. (c) The analytical solution of the heat equation after one time step. (d) The prediction of the IRK-PINN after training. Total L^1 error = 4.81%

Incompressible Navier-Stokes equation: Taylor-Green vortices

The Navier-Stokes Equations describe the motion of Newtonian fluids. In these equations, the variation of the quantity of fluid and its velocity are studied, usually in a compact or periodic domain. The pressure, temperature and density of the fluid also play a role in the description, as well as its viscosity. The search for general solutions to this set of equations is still a very active field of research, with analytical solutions being rarely available for specific assumptions. We study the Taylor-Green vortex [200], which is an incompressible Navier-Stokes equation that has an exact closed form in Cartesian coordinates. The Taylor-Green vortex system is described by

$$\mathcal{N}[u] = u\partial_x u + v\partial_y u + \frac{1}{\rho}\partial_x p - \nu(\partial_x^2 u + \partial_y^2 u) \tag{6.13a}$$

$$\mathcal{N}[v] = u\partial_x v + v\partial_y v + \frac{1}{\rho}\partial_y p - \nu(\partial_x^2 v + \partial_y^2 v), \tag{6.13b}$$

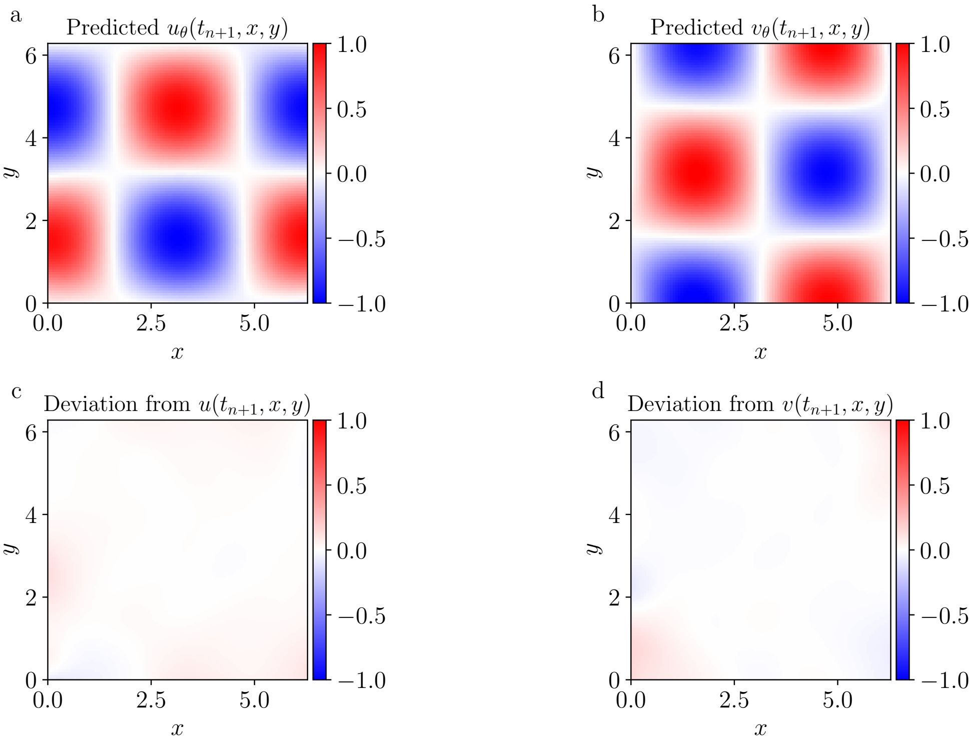

Figure 6.4: Fit to the Taylor-Green vortex data with $\nu = 1$, $\rho = 2$ using a IRK-PINN of order $q = 100$ and time step $\Delta t = 1$. The training was carried out for 20000 epochs with ADAM and 20000 epochs with L-BFGS-B, using an MLP with 10 hidden layers with 16 nodes each and $SiLU$ activation function, and a sampling of $N = 300$ points. (a) and (b) The results of the trained neural network prediction for $u(t_{n+1}, x)$ and $v(t_{n+1}, x)$, respectively. (c) and (d) Deviation from the analytical solution in (6.14). Total L^1-error = 2.01%

where $u(t, x, y)$ and $v(t, x, y)$ are respectively the x and y components of the velocity fields of the fluid, ν is the viscosity, ρ is the mass density, and $p(t, x, y)$ is the pressure of the fluid. To account for the continuity equation for these velocities, which is $\partial_x u + \partial_y v = 0$, we treat it as a boundary condition and add it to our algorithm as an extra term to the loss function.

By restricting the domain of the solutions to the 2-dimensional plane $\mathbf{x} \in \Omega = [0, 2\pi]^2$, the analytical solution takes the form

$$u(t, x, y) = \sin(x)\cos(y)\exp(-2\nu t) \tag{6.14a}$$

$$v(t, x, y) = -\cos(x)\sin(y)\exp(-2\nu t) \tag{6.14b}$$

$$p(t, x, y) = -\frac{\rho}{4}(\cos 2x + \cos 2y)\exp(-4\nu t). \tag{6.14c}$$

For our simulations, we fixed both the viscosity $\nu = 1$ and the mass density $\rho = 2$. In addition, we also fed the pressure $p(t, x, y)$ to the PINN such that its objective is only to predict both components of the velocity fields u and v. The results of the simulation and the details of the neural networked employed are shown in Fig. 6.4. The IRK-PINN is able to learn the solution for a large time step with high accuracy, even when both the input and output dimensions are not one-dimensional.

7 Future outlook

While working on the development of the methodologies discussed in the previous chapters, several novel ideas that hold the potential for further improvement of the standard methods emerged. This chapter provides a description of these ideas, outlining their theoretical motivation, anticipated benefits and possible limitations. Consistent with the general spirit of this thesis, the aim is to deepen the understanding of standard methods for solving the TISE and TDSE, and integrate complementary methodologies for improved performance in high-dimensional problems.

7.1 Enhanced contraction methods for the TISE

In this dissertation, we introduced and developed the concept of enhanced basis sets *via* composition with a normalizing flow, yielding flexible parametrizations of orthonormal basis set, see Chapter 2. We demonstrated that this approach leads to orders-of-magnitude improvements in the accuracy of standard spectral methods, both for a fixed number of functions and for the convergence with respect to the number of functions in the truncated set. The normalizing flows algorithm represents a significant advance over traditional methods by reducing the number of required functions per dimension to achieve similar accuracy.

The reliance on a direct-product basis set composed of univariate functions remains a major computational bottleneck, as this construction does not reduce the arity of the problem. In other words, while less functions per dimension are required in the normalizing-flows approach, the overall scaling of the method remains exponential with the dimensionality. A reduction of arity can be achieved *via* multivariate contraction schemes, as discussed in subsection 1.4.3.

In this section, we propose two different strategies for enhancing the contraction algorithm. First, we explore a deeper understanding of the contraction scheme that motivates the development of a contraction-learning algorithm. Second, we describe a hybrid approach that combines normalizing-flows coordinates with contraction, aiming to exploit the strengths of both methods. Each strategy target different aspects of the contraction algorithm, which provides the possibility of combining strategies.

7.1.1 Learning the contraction scheme

In subsection 1.4.3, we introduced the multivariate contraction scheme and VSCF methods for solving the vibrational TISE in a domain $\Omega \subseteq \mathbb{R}^d$. These techniques avoid the computational bottleneck associated with constructing and diagonalizing the Hamiltonian matrix in a direct-product basis set of univariate functions. Instead, lower-dimensional TISEs are solved and the basis set is built from direct products of the first few solutions.

The outcome of each contraction step is an unitary transformation matrix, obtained from the diagonalization of the reduced-dimensionality TISEs. In fact, these lower-dimensional problems are constructed and solved with the sole purpose of obtaining the unitary transformation that lies an optimal approximation of the full-dimensional eigenfunctions. Nonetheless, the obtained matrix might be far from optimal due to a bad selection of reference geometry and the inability to fully capture multidimensional corrections. Even if the eigenfunctions within each subdomain can be computed with high-precision, no formal guarantees are given to the convergence of the eigenfunctions of the multidimensional system.

To adress this issue, we propose an enhanced contraction scheme that preserves the standard structure of contraction algorithms while incorporating multidimensional information to enhance expressivity. Specifically, for each subdomain, m_i functions denoted as $\left\{\psi_{l_i}(\mathbf{r}^i)\right\}_{l_i=0}^{m_i-1}$ are constructed as parametrized linear combinations of n_i orthonormal basis functions $\left\{\phi_a^i(\mathbf{r}^i)\right\}_{a=0}^{n_i-1}$ defined on the same subdomain as

$$\psi_{l_i}(\mathbf{r}^i;\alpha) = \sum_{a=0}^{n_i} c^i_{l_i a}(\alpha)\ \phi^i_a(\mathbf{r}^i), \tag{7.1}$$

where $m_i \ll n_i$, $\mathbf{c}^i(\alpha)$ is the i-th parametrized unitary matrix that performs the contraction, and α is the parameter set. The number of input functions n_i per subdomain scales exponentially with the dimensionality of the domain, while the output functions m_i does not.

In this step, the expertise for performing the contraction and the computational effort of solving the reduced-dimensionality TISEs are substituted by the computational demand of searching for optimal unitary transformations $\mathbf{c}^i(\alpha)$. The representation of a unitary matrix in terms of a parameter set introduces a new challenge in the construction. A second challenge arises from the method's ability to reliably obtain optimal contraction matrices, avoiding local minima. To enhance stability, we propose initializing the coefficients using the standard contraction coefficients $\mathbf{c}_0^i$ as

$$\mathbf{c}^i(\alpha) = \mathbf{c}_0^i \cdot \mathbf{A}(\alpha), \tag{7.2}$$

where $\mathbf{A}(\alpha)$ is a parametrized unitary matrix and the multiplication of two unitary matrix creates an unitary matrix. It is important to note that this initialization of $\mathbf{c}^i(\alpha)$ is arbitrary; the choice of $\mathbf{c}_0^i$ can be any other unitary matrix such as the identity. However, initializing with the standard contraction matrix and setting $\mathbf{A}$ to the identity provides a favorable starting point for the optimization. If the standard contraction is already optimal, results remain unaffected. If the standard contraction is suboptimal, the optimization will produce enhanced results.

A parametrized unitary matrix $\mathbf{A}(\alpha)$ can be generated using Cayley transform [201]

$$\mathbf{A}(\alpha) = (\mathbb{I} - \mathbf{Q}(\alpha)) \cdot (\mathbb{I} + \mathbf{Q}(\alpha))^{-1}, \tag{7.3}$$

where $\mathbf{Q}(\alpha)$ is a skew-symmetric matrix that depends on the parameter set α, and $\mathbb{I}$ is the identity matrix. The matrix $(\mathbb{I} + \mathbf{Q})$ is invertible by definition of $\mathbf{Q}$.

With this step, the challenge of creating a unitary matrix is mapped to the simpler problem of generating a skew-symmetric parametrizable transformation. The matrix $\mathbf{Q}$ is chosen to have real entries. Let the parameter set be defined by a real-valued matrix α_{ij}. A parameter-dependent real skew-symmetric matrix can be generated as

$$Q_{ij}(\alpha) = \frac{1}{2}(\alpha_{ij} - \alpha_{ji}). \tag{7.4}$$

By construction, $\mathbf{Q}$ is independent of the diagonal entries of the matrix α_{ij}. In addition, the off-diagonal elements of $\mathbf{Q}$ are antisymmetric, further reducing the degrees of freedom of the representation. To minimize the number of parameters in the representation, α_{ij} is created as an upper-triangular matrix with zeros in the diagonal. Given that $\mathbf{Q}$ is a $n_i \times n_i$ matrix, this accounts for a total of $\frac{n_i\ (n_i-1)}{2}$ parameters. We believe that given that only m_i rows of this matrix are used in the contraction, the parameter set has the potential to be reduced even further.

The set of optimised unitary transformation produces a basis set of each subdomain that can be used to build a basis set of the complete domain as a direct product, following standard contraction techniques. The parameters α can be optimised to minimize the M approximated energies of the d-dimensional problem in the contracted basis set. The objective of the optimization is to minimize the loss function

$$\mathcal{L}(\alpha) = \sum_{i=0}^{M-1} \hat{E}_i(\alpha), \tag{7.5}$$

where the energies $\hat{E}_i$ are computed in the contraction basis set.

Since all the produced functions of the contraction are built in close relation to the system, all the approximated energies can be directly incorporated into the optimization. Consequently, the number of functions after contraction N equals the number of approximated energies, $N = M$. This significantly reduces the size of the Hamiltonian matrix that must be computed. Due to the invariance of the trace under unitary transformations, the sum of the M approximated energies can be effeciently computed as the sum of the diagonal elements of the Hamiltonian matrix. As a result, there is no need to store any of the elements of the Hamiltonian matrix in memory during optimization and diagonalization is avoided. Instead, the M diagonal elements diagonal can be added at computation time. This provides a phenomenal relief in the memory costs of the optimization.

The learned contraction is built upon the standard multivariate contraction scheme. The selection of a reference position for the contraction is avoided. The training of the coefficient matrix is capable of learning from the d-dimensional problem to enhance the creation of contracted basis sets for each subdomain, compensating for a possible bad choice of reference geometry. In contrast to standard contraction schemes where the basis set is totally determined by the reference geometry, the optimization of the contraction would arrive to the optimal choice of contracted basis set per subdomain. Assuming that the basis of each subdomain is expressive enough, this method is capable of producing the best possible basis set functions for solving the d-dimensional TISE at hand, under the restriction of a direct prouct basis of functions in each subdomain.

7.1.2 Enhanced contractions *via* normalizing flows

Here, we explore combining the expressivity of the normalizing-flows coordinates with the computational benefits of a multivariate contraction. The applicability and expected results of such an algorithm are explored in this subsection.

Identically to the original normalizing-flow work, the aim is to solve a vibrational TISE defined in a domain $\Omega \subseteq \mathbb{R}^d$ whose coordinates are denoted by the variable $\mathbf{r}$ and has a given Hamiltonian H, defined by its kinetic energy matrix and PES.

In standard multivariate contraction schemes, a partition of the domain $\Omega = \Omega_0 \times \cdots \times \Omega_m$ is chosen, where each domain Ω_i is d_i-dimensional such that $\sum_i d_i = d$. A contraction scheme can be performed for the i-th domain by fixing the coordinates that do not belong to this set to their reference geometry $\mathbf{r}_0$ and solving the TISE for the coordinates of the domain, which we denote by $\mathbf{r}^i$.

Instead of directly performing this contraction, a coordinate transformation is applied. The transformation is parametrised by an invertible neural network $h_\theta : \Omega \to \Gamma$, where

$\mathbf{q} = h_\theta(\mathbf{r})$. This step defines the normalizing-flows coordinate $\mathbf{q}$. As shown in Chapter 2, the normalizing-flows coordinates effectively modify the original TISE by morphing its operators and creates an equivalent TISE defined in the domain Γ.

In the original work, the space Γ was divided in one-dimensional subdomains and the basis set was constructed as direct product of univariate functions of each subdomain. However, by construction, Γ can also be divided into multidimensional subdomains $\Gamma = \Gamma_0 \times \cdots \times \Gamma_m$, where each domain Γ_i is d_i-dimensional such that $\sum_i d_i = d$. While the dimensionality of each subdomain does not need to resemble that of standard contractions, this assumption provides a good starting point.

A contraction can be performed in each of the spaces $L^2(\Gamma_i)$. For solving the i-th TISE, whose variables we denote as $\mathbf{q}^i$, we freeze the rest of the variables to some reference value $\mathbf{q}_0$. In contrast to traditional contraction methods where the reference geometry is chosen according to the PES, e. g., it is the equilibrium position, this reference value is not necesarilly based on any physical intuition. We suggest placing $\mathbf{q}_0$ according to the choice of basis set per dimension. For polynomial-based basis sets, the reference point can be placed at the center of the polynomial expansion. The physical position of the reference geometry now depends on the parameters of the transformation as $\mathbf{r}_0(\theta) = h_\theta^{-1}(\mathbf{q}_0)$. This lowers the expertise barrier to perform the contractions. The optimisation will automatically place $\mathbf{r}_0(\theta)$ in an optimal location.

The basis set of the subdomain Γ_i is created by solving the d_i-dimensional TISE given by

$$\left[\sum_{a,b} \frac{\partial}{\partial \mathbf{q}^i_a} G^{ab}_\theta(\mathbf{q}^i;\mathbf{q}_0) \frac{\partial}{\partial \mathbf{q}^i_b} + V_\theta(\mathbf{q}^i;\mathbf{q}_0)\right] \psi_{l_i}(\mathbf{q}^i;\theta) = E^\theta_{l_i}\, \psi_{l_i}(\mathbf{q}^i;\theta), \tag{7.6}$$

where all quantities are expressed directly in terms of the variable of the problem for simplicity. The kinetic energy matrix G^{ab}_θ and the effective potential V_θ include the required transformations that depend on θ for this equation to be the evaluation of the TISE in the new subset of coordinates. The equation with the necessary transformations in the weak formulation for these operators can be found in equations (2.2), (2.3) and (2.4).

To solve each of these equations, an underlying basis set of the space is required. A basis set $\left\{ \phi^\theta_n \right\}_{n=0}^\infty$ is created using the normalizing flows ansatz, that is

$$\phi_n(\mathbf{q}^i;\theta) = \gamma_n(\mathbf{q}^i)\sqrt{|1/\det J_{h_\theta^{-1}}(\mathbf{q}_i;\mathbf{q}_0)|}, \tag{7.7}$$

where $\gamma_n(\mathbf{q}^i)$ is the underlying direct-product orthonormal basis set of $L^2(\Gamma_i)$ and the determinant is calculated at the reference point for the coordinates that are not in the set.

Identically to standard multivariate contraction schemes, solving Equation (7.6) using a total of n_i of such functions creates a set of m_i orthogonal functions $\left\{ \psi^\theta_{l_i} \right\}_{l_i=0}^{m_i}$, $m_i \ll n_i$.

As a final step, the complete TISE is solved in the span of direct products of the functions obtained for each contraction step. The Hamiltonian matrix is computed in the representation of the d-dimensional direct-product functions and is diagonalised to obtain the approximated energies. An optimization process for the normalizing-flows parameters can be set up by choosing the number of functions per dimension

n_i, the number of contracted functions per subdomain m_i, and by minimizing the M approximated energies that result of the contraction basis set.

The normalizing-flows coordinates morph the operators in the TISE and enhance their mode separability. This enhanced separability improves the contraction scheme in two different ways. Within each subdomain, the convergence of the eigenfunctions of the TISE is improved by the enhanced separability of the operators. An example of this separability is ilustrated in Figure 2.4 for the HCN molecule. While this allows for less expensive numerical calculations in each subdomain, contraction assumes the abilty to converge the first m_i eigenfunctions (as many functions can be used in each subdomain). The enhancement in computations in each subdomain only results in significant improvements when this assumption does not hold true.

On the other hand, normalizing flows facilitate a greater degree of decoupling between coordinates belonging to different partitions of the domain. If the operators are completely separable in each subdomain, the direct product of contractions produces the set of eigenfunctions directly. This adjustment of the system to the direct product ansatz is not attainable through conventional spectral methods. Consequently, normalizing-flows-based contractions would result in improvements of orders of magnitude in accuracy with respect to standard contraction for the same number of contracted functions.

To further reduce the numerical costs of solving the TISE in high-dimensional systems, the normalizing-flows contraction can be combined with the contraction-learning scheme introduced in the previous subsection. In this combined approach, the optimal coordinate transformation is learned by the normalizing flow algorithm, while the contraction in the transformed coordinates is simultaneously optimized.

As a result, the repeated solution of the TISEs within each subdomain is replaced by the optimization of the entries of the unitary matrix. This enables the direct construction of the contracted functions for the d-dimensional problem, eliminating the need of subsequent diagonalizations in the training loop. This significantly reduces the computational overhead associated to tracking the normalizing-flow operations through repeated diagonalizations.

This integrated framework also enables an efficient two-step contraction strategy. In the first step, the normalizing-flows is trained to optimize the separability of operators, potentially using minimal or no information of the basis set per subdomain. When a good coordinate set is found, the second step involves learning the parameters of the contraction in the optimized coordinate representation. This two-step method combines an optimal subdomain partition, provided by the normalizing flows change of coordinates, with the optimal basis set per subdomain, obtained by the learned contraction. Although the technical implementation is challenging and demans sophisticated software engineering, this construction has the potential to evolve into a state-of-the-art strategy for solving the TISE in high-dimensional systems.

7.2 Enhancing solutions of the TDSE using parameter-dependent basis sets

As discussed in subsections 1.3.3 and 1.3.4, a variety of methods have been explored in the literature for solving the TDSE. Different methods yield distinct accuracies, due to the bigger or smaller flexibility of the ansatze. The quality of the simulation for each method also depends largely on the complexity of the dynamics. A small perturbation to the static Hamiltonian produces dynamics that can be easily resolved using spectral methods in the eigenbasis, while it can be as challenging as more intense perturbations for PINNs. In contrast, an intese perturbation creates dynamics that require many basis functions for its resolution, deteriorating the prediction quality of spectral methods.

Spectral methods are the standard choice for predicting molecular dynamics, given that spatial basis sets guarantee convergence to the solution at each time. In contrast to solving the TISE where the variational limit is used as a metric of the quality of the numerical results, there does not exist a metric that clearly shows the quality of the TDSE predictions. A common recipe to measure the quality of a prediction is computing the dynamics using two slightly different basis set size truncations. If both predictions yield similar results, the calculation is accepted as a converged simulation.

While this is often a good indicator of the quality of the approximation, it does not provide any rigurous guarantees. For example, if a state that is important for the dynamics is not contained in neither of the truncations, both simulations may produce similar results while being unable to describe the true dynamics of the system.

In this section, we use the flexibility of the parametrizable basis sets created in the normalizing-flows coordinates to enhance the computational capabilities of spectral methods in approximating solutions of the TDSE.

7.2.1 Optimal static basis set *via* normalizing flows

The main goal of this subsection is to propose a metric to measure the quality of the basis sets in TDSE spectral calculations, such that the quality of the approximation can be estimated *a priori*. Based on this metric, we propose an algorithm that is capable of improving the accuracy of TDSE simulations for a given basis set truncation.

The spatial component of the wavefunction is contained in the Hilbert space defined on the domain Ω, which is characteristic to the problem. Any function that belongs to this space can be written in terms of a linear combination of the functions in a basis set $\{\phi_i\}_{i=0}^{\infty}$. Due to the finite memory available for computation, only N functions of the basis set can be used for performing the approximation. This set of functions does not span the complete Hilbert space, but a subdomain of it. We denote the subspace of the first N functions of the basis set by $L^2_N(\Omega)$ and define it as the representation space. In other words,

$$L^2_N(\Omega) = \left\{ f \in L^2(\Omega) \mid f = \sum_{i=0}^{N-1} c_i\phi_i, \quad c_i \in \mathbb{C} \right\}. \tag{7.8}$$

From this point on, we assume the implicit domain of all functions is Ω and avoid its repetition.

As disccussed in subsection 1.3.3, the time-dependency of the TDSE can be encoded in the linear coefficients of the expansion. In the representation subspace, the wavefunction is approximated by

$$|\Psi(t)\rangle \approx \sum_{i=0}^{N-1} c_i(t)|\phi_i\rangle. \tag{7.9}$$

The standard spectral method solution is obtained by solving the equations of motion of $c_i(t)$ defined in (1.36). An analytical solution of the EOM produces an approximation that is *optimal for the representation space*. However, the analyticity of the approximation does not guarantee the exactness of the solution to the TDSE. While this statement may seem contradictory, it is precisely the fact that the representation space is constrained to a subspace of the Hilbert space that may lead to inaccurate approximations.

This problem can be ilustrated by an example. Let two different basis sets that are identically truncated at N functions have representation subspaces denoted by $L^2_{N,(0)}$ and $L^2_{N,(1)}$. Solve the EOM within the different subspaces. Which of the two offers a better solution to the TDSE? Each basis set produces a different results, but in both cases the EOM can be solved exactly within the representation.

The metric of the quality of the solution obtainable in a representation space is not contained in said space. Instead, this information is contained in its orthonormal complement, i. e., the subspace of functions that belong to the Hilbert space but are not contained in L^2_N. The orthonormal complement $\hat{L}^2_N$ is defined as

$$\hat{L}^2_N = \left\{ f \in L^2(\Omega) \mid \langle f|g\rangle = 0, \quad \text{for all } g \in L^2_N \right\} \tag{7.10}$$

L^2_N is a good subspace for simulating the dynamics of a TDSE if the projection of the true solution to $\hat{L}^2_N$ is negligible, while it fails capturing the dynamics if it is significant at any time.

Let the set $\left\{\hat{\phi}_i\right\}_{i=0}^{\infty}$ be an orthonormal basis set of L^2. In particular, if L^2_N was built truncating the functions of a basis set, $\hat{L}^2_N$ is the set of basis functions $\left\{\phi_i\right\}_{i=N}^{\infty}$. If the approximation in the projected space solves the TDSE exactly, the evolution of the approximated wavefunction into $\hat{L}^2_N$ for all subsequent time steps is zero. Assume that the initial condition of the TDSE is expressable in L^2_N. Then, the TDSE is solved exactly within the representation space if the dynamics satisfy

$$\int_{t_0}^{T} dt \; \langle \hat{\phi}_i | \partial_t \Psi(t) \rangle = \int_{t_0}^{T} dt \; \sum_{j=0}^{N-1} c_j(t) \langle \hat{\phi}_i | H(t) | \phi_j \rangle = 0, \quad \forall i, \tag{7.11}$$

where the representation of the wavefunction in the functions of the basis of equation (7.9) was introduced. By construction of the solution of the equations of motions in L^2_N, this condition is systematically omitted.

Equation (7.11) provides a metric of the quality of the approximation. Increasingly bigger deviations from this equation yield increasingly worse approximations.

The main component of equation (7.11) are the terms of the Hamiltonian matrix that couple the different subspaces $\langle \hat{\phi}_i | H(t) | \phi_j \rangle$. If these terms are zero at all times, the equation is always satisfied, independently of the values of $c_j(t)$. We define this matrix as the Hamiltonian off-diagonal block— following the standard notation that the diagonal block of a matrix connects a subspace with itself.

Assume that the Hamiltonian can be separated on a static component H_0 and a time-dependent component $h(t)$. In molecular physics, these are the Hamiltonian of the molecule and the of the perturbation respectively. In addition, assume the perturbation Hamiltonian is separable into a spatial and a temporal component, i. e. $h(t) = X(\mathbf{x})\,T(t)$. This assumption is valid in perturbations by external fields with a wavelength that is bigger than the scale of the system. In this case, the molecule effectively percieves the same electric field at all points, making the spatial and time components separable.

Under this division, the off-diagonal block of the Hamiltonian can be written as

$$\langle \hat{\phi}_i | H(t) | \phi_j \rangle = \langle \hat{\phi}_i | H_0(\mathbf{x}) | \phi_j \rangle + T(t) \langle \hat{\phi}_i | X(\mathbf{x}) | \phi_j \rangle. \tag{7.12}$$

The representation must satisfy $\langle\hat{\phi}_i|H_0(\mathbf{x})|\phi_j\rangle = 0$ and $\langle\hat{\phi}_i|X(\mathbf{x})|\phi_j\rangle = 0$ to produce an off-diagonal block that is zero at all times. While the first of these equation is satisfied in the set of eigenfunctions of H_0, the second one may not.

While these concepts are not novel, their domain of application was narrow. For static basis sets, checking this a priori measure can be mapped to running calculations on two different basis set truncations. However, the recent advances in the creation of flexible parametrized basis sets allow for the optimisation of the representation space. This is the case of the normalizing flows methodology. We provide a description of this method.

Using the normalizing-flows ansatz (equation (2.1)), we denote by $\left\{\phi_i^\theta\right\}_{i=0}^{N-1}$ the functions in the representation space, where θ refers to the parameters of the invertible neural network. As the projected space is now parameter-dependent, we denote it as $L_N^2(\theta)$. Identically, the orthonormal complement is denoted as $\hat{L}_N^2(\theta)$.

Due to the limis in computational memory, M representatives of the complementary space are chosen to satisfy equation (7.12). These functions are simply the next M functions of the basis set that do not enter the representation, that is $\left\{\hat{\phi}_i^\theta\right\}_{i=0}^{M-1} = \left\{\phi_j^\theta\right\}_{j=N}^{N+M-1}$. Assume that, if the contribution of this set of functions to the solution is negligible, the contribution of the other functions in the infinite set $\hat{L}_{N,\theta}^2$ is also negligible.

Under this parametric representation, equation (7.11) can be used as a metric for optimisation. A set of optimal parameters for the time-dependent solution of the TDSE are those who minimise the loss function

$$\mathcal{L}_H = \sum_{i=0}^{N-1}\sum_{j=0}^{M-1}\int_{t_0}^{T} dt' \left|\langle\hat{\phi}_j^\theta|H(t')|\phi_i^\theta\rangle\right|^2 . \tag{7.13}$$

Under the assumption that the Hamiltonian is the sum of the static and time-dependent perturbation components, and that the perturbation is separable in spatial and time components, the loss function can be transformed to

$$\mathcal{L}_H = \sum_{i=0}^{N-1}\sum_{j=0}^{M-1} \left|\langle\hat{\phi}_j^\theta|H_0(x)|\phi_i^\theta\rangle\right|^2 + \epsilon\left|\langle\hat{\phi}_j^\theta|X(x)|\phi_i^\theta\rangle\right|^2 , \tag{7.14}$$

where $\epsilon > 0$ is a constant that weights the relative importance of each term of the loss function. ϵ can be chosen as the maximum value of $T(t)$ to ensure that this contribution is minimal even at its peak.

An extra term to the loss function is required to ensure that the projection to the initial state of the propagation $|\psi_0\rangle$ is accurate. This is an important requirement, as it is the initial condition of the propagation together with the perturbation Hamiltonian that determine the dynamics. If the information of the initial state is available, the projection onto the representation space can be used as a loss function. The approximation $|\psi_0\rangle \approx \sum_{i=0}^{N-1} c_j|\phi_j^\theta\rangle$ is exact if $\sum_{i=0}^{N-1}|c_j|^2 = 1$. A loss function for the initial condition can be created as

$$\mathcal{L}_{init}(\theta) = -\sum_{i=0}^{N-1}\left\|\langle\phi_j^\theta|\psi_0\rangle\right\|^2 . \tag{7.15}$$

If the time-dependent dynamics start from the unperturbed molecule in the ground state, no prior knowledge is required for the loss function of the initial conditions. Instead, the energy of the static Hamiltonian H_0 can be minimized within the basis representation. This is equivalent to optimizing to the variational limit. The loss function for the initialization can be then written as

$$\mathcal{L}_{init}(\theta) = \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} c_i^* c_j \langle \phi_i^\theta | H_0 | \phi_j^\theta \rangle, \tag{7.16}$$

where c_j are the coefficients of the projection to the ground state obtained by diagonalization of the static Hamiltonian matrix. When this condition is properly satisfied, the initial condition is created by setting $c_j(t_0) = c_j$.

For the complete optimization of the projected space, the loss function is composed of the weighted sum of the block-diagonal loss function in (7.14) and the initial condition loss function in (7.16).

Failing to minimize (7.14) might still yield in aceptable solutions of the TDSE, due to having a non-zero Hamiltonian element between a state that does not belong to the representation and a state that is in the projected space but does not have any significant contribution to the dynamics. Similarly, if the time propagation spans a short time interval, the dynamics might be well represented within few states. In this case, a small representation space would satisfy the demands of the dynamics. While it is possible that the constraints of the optimization are not always required, the formal guarantee of obtaining a solution of the TDSE outweights these potential inconveniences.

7.2.2 Evolving basis sets using normalizing flows for molecular dynamics

A fundamental difference between the Schrödinger equations is that the TISE aims to calculate as many eigenfunctions as possible. In contrast, for each choice of initial conditions and Hamiltonian, the solution of the TDSE is unique. As long as the functions used in the representation of the wavefunction properly mimic the time-dependency of the problem, not many functions are actually required to obtain an accurate solution. Purely-numerical approaches like PINNs use a continous representation with a single function parametrised by a neural network and try to fit the parameters to solve the time propagation. Other methods aliviate the complexity of fitting a single function by using a set of functions that can be optimized to solve the TDSE, e. g., the MCTDH or gaussian-MCTDH. In this section, we introduce a time-evolving normalizing flows approach for solving the TDSE in the span of few functions of a time-dependent basis set.

The TDSE can be reformulated as the variational Dirac-Frenkel functional equation [202–204]

$$\mathrm{F}[\Psi] = \int_{t_0}^{T} dt \, \left\| i\partial_t \Psi - H\Psi \right\|^2 = \int_{t_0}^{T} dt \, \langle i\partial_t \Psi - H\Psi | i\partial_t \Psi - H\Psi \rangle. \tag{7.17}$$

If the solution to the TDSE, denoted as Ψ_s, is inputed in the variational Dirac-Frenkel equation, both bra and ket will be evaluated to zero making $\mathrm{F}[\Psi_s] = 0$. Given that the solution of the TDSE is unique, $F = 0$ can only be achieved by Ψ_s. As the definition uses the integral of a norm, $F > 0$ for any other function. Therefore, this functional equation is a lower bound that can be used for optimization. As a minimum of F, the solution of the TDSE satisfies the variational equation

$$\left.\frac{\delta \mathrm{F}[\Psi]}{\delta \Psi}\right|_{\Psi=\Psi_s} = 0. \tag{7.18}$$

By applying the variational derivative on the bra and the ket of (7.17), two terms are produced, each being the complex-conjugation of the other. For the real and imaginary part of the equation to be zero, we require both to be zero separately. Then, satisfying equation (7.18) requires

$$\langle \delta\Psi | i\partial_t \Psi H \Psi \rangle = 0, \tag{7.19}$$

where $\delta\Psi$ is any variation of the wavefunction.

We use the normalizing-flows algorithm as the choice of spectral method. Then, the wavefunction is approximated as a linear combination of the normalizing-flows augmented basis set $\left\{\phi_i(\mathbf{r};\theta)\right\}_{i=0}^{\infty}$, where θ are the parameters of the invertible neural network. To augment the expressivity of the basis set over time, we allow the parameters of the normalizing-flows coordinates to be time-dependent $\theta(t)$ and aim to learn their time-evolution— together with the time-evolution of the linear coefficients. The basis set functions are allowed to adjust to the evolution of the molecular dynamics, changing the focus of their resolution power over space in time. Alternatively, a normalizing-flow that accepts time directly as a variable can be used. In this construction, the parameters would be time-independent and the change of variables would evolve in time.

We perform the derivations in this subsection under the ansatz of the time-dependent parameter set. Then, the wavefunction is approximated in the span of N functions of the basis as

$$\Psi = \sum_{i=0}^{N-1} c_i(t)\, \phi_i(\mathbf{r};\theta(t)). \tag{7.20}$$

We explore variations on $c_i(t)$ and $\theta(t)$. The total variation of the wavefunction is the sum of the variations of the individual terms

$$\delta\Psi = \delta c_i \cdot \frac{\delta\Psi}{\delta c_i} + \delta\theta_l \cdot \frac{\delta\Psi}{\delta\theta_l}, \tag{7.21}$$

where θ_l is the l-th parameter of the parameter set. To satisfy equation (7.18) under any possible variation, each contribution to the variation needs to be identically zero. This condition results in the set of equations of motion

$$\delta c_i \langle \frac{\delta\Psi}{\delta c_i} | i\partial_t \Psi H \Psi \rangle = 0, \tag{7.22}$$

$$\delta\theta_l \langle \frac{\delta\Psi}{\delta\theta_l} | i\partial_t \Psi H \Psi \rangle = 0. \tag{7.23}$$

Note that δc_i and $\delta\theta_l$ are non-zero constants and can be disregarded in the overall equation.

By performing the functional and time derivatives of the first equation of motion, we obtain

$$\langle \phi_i \mid \sum_{j=0}^{N-1} i\frac{\partial c_j}{\partial t}\phi_j + ic_j \sum_l \frac{\partial \theta_l}{\partial t}\frac{\partial \phi_j}{\partial \theta_l} - Hc_j\phi_j\rangle = 0. \tag{7.24}$$

Using the orthonormality of the basis functions, the first equation of motion becomes

$$i\frac{\partial c_i}{\partial t} = \sum_{j=0}^{N-1} c_j\langle\phi_i|H\phi_j\rangle - ic_j\sum_l \frac{\partial \theta_l}{\partial t}\langle\phi_i \mid \frac{\partial \phi_j}{\partial \theta_l}\rangle, \tag{7.25}$$

where the objects that are not space-dependent were moved outside the inner-product for visual clarity. By setting $\frac{\partial \theta_l}{\partial t} = 0$, the original EOM for spectral methods (1.36) is recovered, which shows the consistency of this derivation with the previous theory.

The functional and time derivatives can be performed on the EOM for the neural-network parameters, obtaining

$$\sum_{i=0}^{N-1}\langle c_i\frac{\partial \phi_i}{\partial \theta_l} \mid \sum_{j=0}^{N-1} i\frac{\partial c_j}{\partial t}\phi_j + ic_j\sum_m \frac{\partial \theta_m}{\partial t}\frac{\partial \phi_j}{\partial \theta_m} - Hc_j\phi_j\rangle = 0.$$

By rearranging terms, moving the objects that are not spatial-dependent outside of the inner-products, and substituting for the time-derivative of θ_l, the resulting equation of motion can be written as

$$\frac{\partial \theta_m}{\partial t} = \sum_l \left(\mathcal{H}_l - C_l\right) A^{-1}_{lm}, \tag{7.26}$$

where the quantities $A_{lm}, \mathcal{H}_l$ and C_l are defined as

$$\begin{aligned}
A_{lm} &= \sum_{i=0}^{N-1}\sum_{j=0}^{N-1} ic_i^*\, c_j\langle\frac{\partial \phi_i}{\partial \theta_l}\,|\frac{\partial \phi_j}{\partial \theta_m}\rangle \\
\mathcal{H}_l &= \sum_{i=0}^{N-1}\sum_{j=0}^{N-1} c_i^*\, c_j\langle\frac{\partial \phi_i}{\partial \theta_l}|H\phi_j\rangle \\
C_l &= \sum_{i=0}^{N-1}\sum_{j=0}^{N-1} ic_i^*\frac{\partial c_j}{\partial t}\langle\frac{\partial \phi_i}{\partial \theta_l} \mid \phi_j\rangle.
\end{aligned} \tag{7.27}$$

The set of coupled differential EOM, (7.25) and (7.26), cannot be solved analytically. To propagate the quantities in time, the use of numerical methods is necessary. We propose the use of some standard numerical solver, such as implicit or explicit Runge-Kutta methods, to propagate this equation over time.

This approach is similar to g-MCTDH, where the functions of the set are gaussians that can change their position and shape during time, allowing each of them to model a wavepacket. The proposed method pushes this approach even further, by introducing a parametrized time-dependent orthonormal basis set. Here, the functions are orthonormal at all times— avoiding the computation of the inverse of the overlap matrix for normalization.

We predict time-dependent normalizing-flows method to provide an enhanced description of the wavefunction, even when a small basis set is used. At each time, the objective of the optimization is to create a basis set that contains within its linear span the spatial part of the wavefunction at said time. For every given time, this is similar to the approximation problem of finding an eigenstate of the TISE. In chapter 2, it has been shown that the normalizing-flows algorithm is capable of approximating functions with orders of magnitude enhanced precision compared to standard spectral methods. By extrapolating these results to each time step, the time-dependent normalizing flows holds the potential to create highly-accurate approximations of the TDSE.

8 Conclussions

This dissertation focused on the central topic of approximating solutions of the Schrödinger equations, both in the time-independent and time-dependent formalisms. Different numerical methods were used to build efficient calculations that are competitive with state-of-the-art methods. In particular, the presented methodology focus on aliviating the problems related to having high dimensionality.

For the time-independent Schrödinger equation, we improved the convergence of many solutions by the introduction of the normalizing-flows algorithm. This method creates parametrical families of orthonormal basis sets of the Hilbert space by introducing a non-singular change of coordinates. This parametrization allows for optimization of the basis to the problem at hand. When comparing to modern methodologies, the normalizing-flows algorithm is capable of creating more accurate approximations for a bigger set of target states. The enhancement in approximation capabilities was demonstrated in calculations of semi-rigid molecules of different sizes (H_2S and H_2CO), and floppy molecules (the HCN/CNH isomers).

The optimal change of variables depends on the molecule, underlying basis set and basis set truncation. However, we demonstrated that the variance with respect to the basis set truncation becomes small as the number of target states and basis functions increases. This provides the possibility of creating an efficient optimization algorithm: optimize for medium size truncations and then transfer the coordinate to a truncation where optimization would not be computationally feasible. Moreover, the optimal coordinate set can also be transferred across molecules with similar molecular motifs. Under the choice of a product basis of univariate functions, the assignment of quantum numbers in the optimal coordinates is also signficantly improved with respect to the assignment in the original coordinate set. This illustrates the decoupling of the normal modes that enhances the approximation capabilities.

Nonetheless, some challenges remain open regarding the application of normalizing flows in high-dimensional systems. First, we lack an approach to impose periodicity and symmetry in the optimised coordinate set. A succesful implementation of these constraints may allow for a reduction in the neural network complexity, and the production of block-diagonal structures in the Hamiltonian matrix, reducing the computational costs for each block. On the numerical side, the use of quadratures for integrating presents a soft bottleneck for the calculations. In other words, calculations using any number of points can be performed but the use of large quadrature sets considerably hinders the efficiency. To alleviate this bottleneck, we are exploring the implementation of an optimized Monte Carlo importance sampling algorithm in which the number of required points do not scale exponentially with the dimensionality. The application of normalizing flows has been so far restricted to the use of direct products of univariate basis functions. The methodology can be extended by the use of alternative basis sets constructions that do not scale exponentially with the dimensionality of the problem. A posibility is the use of radial basis functions centered at different nodes, where the position of the nodes can be optimized by the coordinate transformation. However, these non-standard basis sets lack of quadratures that are specifically built for them. For this reason, Monte Carlo integration is the sensible choice. Pointwise solutions of the TISE can alternatively be approximated using these sets.

The same difficulties regarding dimensionality in the solution of the TISE are found in other high-dimensional approximation problems. To this extent, we developed a theory to enhance the approximation of continuous functions using polynomials composed with strictly monotonic functions, once more acting as non-singular change of coordinates. We prove that the induced set from this composition remains dense in the space of continuous functions, meaning that any function can be approximated to arbitrary precision in the linear span of the set. We also provide a finite bound for the degree of the polynomial that can achieve arbitrary precision in the approximation of any continuous function. This approach can also be applied to the Schrödinger equation in multidimensional problems. Concretely, we demonstrate that this induced set is capable of creating better approximations of the potential energy surface by applying it to the H_2S molecule.

In this dissertation, we also tackled the challenge of solving the time-dependent problems for complex experiments, concretely, for the laser-induced electron diffraction imaging technique. The use of the semiclassical model allows for the production of simulations of the dynamics that are not feasible by solving the TDSE, due to both the complexity of the dynamics and the dimensionality of the complete problem. While different approximations are introduced, the obtained simulations produce meaningful results that help us understand the role of the molecular geometry in the photoangular distribution. A natural continuation of this project is reducing the number and extent of the approximations. For example, introducing time-dependent molecular electric fields and improved tunneling ionization probability might result in significant improvements on the simulation power. The next keystone for the development of a complete theory of LIED is solving the inverse problem. At the moment, an approximated result of the inverse problem is obtained using reconstruction techniques of ultrafast electron dynamics, that do not account for the many possible rescattering events.

Remarks

Nomenclature

For simplicity, the thesis is written in atomic units, i. e. $\hbar = m_e = e = 1$. $i = \sqrt{-1}$ is used as the imaginary unit.

Abreviations

This is a list of the abreviations used throughout the thesis

- *PDE*: Partial differential equation
- *TDSE*: Time-dependent Schrödinger equation.
- *TISE*: Time-independent Schrödinger equation.
- *PES*: Potential-energy surface.
- *COM*: Center of mass.
- *VSCF*: Vibrational self-consistent-field theory
- *MCTDH*: Multiconfigurational time-dependent Hartree method
- *EOM*: Equation of motion
- *PINN*: Physics-informed neural networks
- *iResNet*: Invertible residual neural network
- *MLP*: Multilayer perceptron
- *iso-HO*: Isotropic harmonic oscillator
- *LIED*: Laser-induced electron diffraction
- *MO-ADK*: Molecular-orbital applied Ammosov-Delone-Krainov theory
- *QRS*: Quantitative rescattering theory
- *IAM*: Independent atom model
- *DCS*: Differential cross section
- *HOMO*: Highest occupied molecular orbital
- *PAD*: Photoangular distribution
- *MCF*: Molecular contrast factor
- *ESP*: Electrostatic potential field
- *RGI*: Regular grid interpolator
- *RMSE*: Root mean square error

- *NN*: Neural network
- *VMI*: Velocity map imaging spectrometer
- *(I)RK*: (Implicit) Runge-Kutta

References

[1] E. Schrödinger, Phys. Rev. **28**, 1049 (1926).

[2] C. Lubich, *From quantum to classical molecular dynamics: reduced models and numerical analysis* (European Mathematical Society, 2008).

[3] L. D. Landau and E. M. Lifshitz, *Quantum Mechanics – Nonrelativistic Theory* (Pergamon Press, Oxford, 1977), 3rd ed., translated from Russian by J.B. Sykes and J.S. Bell.

[4] E. Noether, Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse pp. 235–257 (1918), URL `http://eudml.org/doc/59024`.

[5] M. Porta, *Mathematical quantum theory* (2019), lecture notes (accessed 2022-03-15), URL `https://www.math.uni-tuebingen.de/de/forschung/maphy/lehre/ws-2019-20/mqt/notes.pdf`.

[6] M. Born and R. Oppenheimer, Ann. Physik **84**, 457 (1927), URL `https://doi.org/10.1002/andp.19273892002`.

[7] H. H. Nielsen, Rev. Mod. Phys. **23**, 90 (1951), URL `https://link.aps.org/doi/10.1103/RevModPhys.23.90`.

[8] G. Herzberg, *Molecular Spectra and Molecular Structure: Electronic Spectra and Electronic Structure of Polyatomic Molecules*, vol. 3 (Krieger Publishing Company, Malabar, Fl, USA, 1991).

[9] D. J. Griffiths and D. F. Schroeter, *Introduction to Quantum Mechanics* (Cambridge University Press, 2018), 3rd ed.

[10] P. R. Bunker and P. Jensen, *Molecular Symmetry and Spectroscopy* (NRC Research Press, Ottawa, Ontario, Canada, 1998), 2nd ed.

[11] B. Podolsky, Phys. Rev. **32**, 812 (1928), URL `https://link.aps.org/doi/10.1103/PhysRev.32.812`.

[12] M. Reed and B. Simon, *Methods of Modern Mathematical Physics: Functional analysis*, Methods of Modern Mathematical Physics (Academic Press, 1980), ISBN 9780125850506, URL `https://books.google.de/books?id=bvuRuwuFBWwC`.

[13] R. d. l. Madrid, European Journal of Physics **26**, 287–312 (2005), URL `http://dx.doi.org/10.1088/0143-0807/26/2/008`.

[14] M. Feit, J. Fleck, and A. Steiger, J. Comput. Phys. **47**, 412 (1982), URL `https://doi.org/10.1016/0021-9991(82)90091-2`.

[15] H.-D. Meyer, U. Manthe, and L. Cederbaum, Chem. Phys. Lett. **165**, 73 (1990), ISSN 0009-2614, URL `http://www.sciencedirect.com/science/article/pii/000926149087014I`.

[16] H.-D. Meyer, F. Gatti, and G. A. Worth, eds., *Multidimensional Quantum Dynamics: MCTDH Theory and Applications* (Wiley, 2009), URL `https://doi.org/10.1002/9783527627400`.

[17] G. A. Worth, H.-D. Meyer, H. Köppel, L. S. Cederbaum, and I. B. and, International Reviews in Physical Chemistry **27**, 569 (2008), URL `https://doi.org/10.1080/01442350802137656`.

[18] M. Raissi, P. Perdikaris, and G. Karniadakis, J. Comput. Phys. **378**, 686 (2019), URL `https://doi.org/10.1016/j.jcp.2018.10.045`.

[19] R. Kadison and J. Ringrose, *Fundamentals of the Theory of Operator Algebras. Volume I*, Fundamentals of the Theory of Operator Algebras (Adv. Mass Spectrom., 1997), ISBN 9780821808191, URL `https://books.google.de/books?id=Q3J6TV6euVYC`.

[20] G. Szegő, *Orthogonal Polynomials*, American Mathematical Society colloquium publications (American mathematical society, 1939), URL `https://books.google.de/books?id=AxsPAAAAIAAJ`.

[21] S. Carter and N. Handy, Computer Physics Communications **51**, 49 (1988), ISSN 0010-4655, URL `https://www.sciencedirect.com/science/article/pii/0010465588900616`.

[22] X.-G. Wang and T. Carrington, J. Chem. Phys. **117**, 6923 (2002), URL `https://doi.org/10.1063/1.1506911`.

[23] X.-G. Wang and T. Carrington, J. Chem. Phys. **129**, 234102 (2008), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/1.3027825`.

[24] J. Koput, S. Carter, and N. C. Handy, The Journal of Chemical Physics **115**, 8345 (2001), ISSN 0021-9606, URL `https://doi.org/10.1063/1.1410976`.

[25] J. M. Bowman, Acc. Chem. Res. **19**, 202 (1986), URL `https://doi.org/10.1021/ar00127a002`.

[26] J. M. Bowman, The Journal of Chemical Physics **68**, 608 (1978), ISSN 0021-9606, URL `https://doi.org/10.1063/1.435782`.

[27] Z. Bačić, R. B. Gerber, and M. A. Ratner, J. Chem. Phys. **90**, 3606 (1986), URL `https://doi.org/10.1021/j100407a029`.

[28] P. Davis and P. Rabinowitz, *Methods of Numerical Integration*, Computer Science and Applied Mathematics (Academic Press, 1984), ISBN 978-0-12-206360-2, URL `https://doi.org/10.1016/C2013-0-10566-1`.

[29] J. Dagpunar, *Simulation and Monte Carlo: With Applications in Finance and MCMC* (John Wiley & Sons, Ltd, 2007), ISBN 9780470854945.

[30] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods* (Springer New York, NY, 2004), 2nd ed.

[31] Z. Li, Y.-l. Fu, Z. Luo, S. Yang, Y. Wu, H. Wu, G. Wu, W. Zhang, B. Fu, K. Yuan, et al., Science **383**, 746–750 (2024), ISSN 1095-9203, URL `http://dx.doi.org/10.1126/science.adn3357`.

[32] D. J. Auerbach, J. C. Tully, and A. M. Wodtke, Nat. Sci. **1**, e10005 (2021), ISSN 2698-6248, URL `http://dx.doi.org/10.1002/ntls.10005`.

[33] C. D. Foley, C. Xie, H. Guo, and A. G. Suits, Science **374**, 1122–1127 (2021), ISSN 1095-9203, URL `http://dx.doi.org/10.1126/science.abk0634`.

[34] B. Margulis, K. P. Horn, D. M. Reich, M. Upadhyay, N. Kahn, A. Christianen, A. van der Avoird, G. C. Groenenboom, M. Meuwly, C. P. Koch, et al., Science **380**, 77–81 (2023), ISSN 1095-9203, URL `http://dx.doi.org/10.1126/science.adf9888`.

[35] I. Rahinov, A. Kandratsenka, T. Schäfer, P. Shirhatti, K. Golibrzuch, and A. M. Wodtke, Phys. Chem. Chem. Phys. **26**, 15090–15114 (2024), ISSN 1463-9084, URL `http://dx.doi.org/10.1039/d4cp00957f`.

[36] G. Meng, C. Hu, and B. Jiang, J. Phys. Chem. C **126**, 12003–12008 (2022), ISSN 1932-7455, URL `http://dx.doi.org/10.1021/acs.jpcc.2c03056`.

[37] Y. Yoneda and H. Kuramochi, J. Phys. Chem. A **127**, 5276–5286 (2023), ISSN 1520-5215, URL `http://dx.doi.org/10.1021/acs.jpca.3c02489`.

[38] S. O. M. Wright, I. Waldmann, and S. N. Yurchenko, Mon. Not. R. Astron. Soc. **512**, 2911–2924 (2022), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stac654`.

[39] S. O. M. Wright, S. K. Nugroho, M. Brogi, N. P. Gibson, E. J. W. de Mooij, I. Waldmann, J. Tennyson, H. Kawahara, M. Kuzuhara, T. Hirano, et al., Astronomical J. **166**, 41 (2023), ISSN 1538-3881, URL `http://dx.doi.org/10.3847/1538-3881/acdb75`.

[40] N. H. Pinkowski, Y. Ding, C. L. Strand, R. K. Hanson, R. Horvath, and M. Geiser, Meas. Sci. Technol. **31**, 055501 (2020), URL `https://dx.doi.org/10.1088/1361-6501/ab6ecc`.

[41] A. Ehn, J. Zhu, X. Li, and J. Kiefer, Appl. Spectrosc. **71**, 341 (2017), pMID: 28155328, URL `https://doi.org/10.1177/0003702817690161`.

[42] S. N. Yurchenko, W. Thiel, and P. Jensen, J. Mol. Spectrosc. **245**, 126 (2007), URL `https://doi.org/10.1016/j.jms.2007.07.009`.

[43] E. Mátyus, G. Czakó, and A. G. Császár, J. Chem. Phys. **130**, 134112 (2009), URL `https://doi.org/10.1063/1.3076742`.

[44] J. M. Bowman, S. Carter, and X. Huang, Int. Rev. Phys. Chem. **22**, 533 (2003), URL `https://doi.org/10.1080/0144235031000124163`.

[45] O. Christiansen, J. Chem. Phys. **120**, 2140–2148 (2004), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/1.1637578`.

[46] M. B. Hansen, M. Sparta, P. Seidler, D. Toffoli, and O. Christiansen, J. Chem. Theory Comput. **6**, 235–248 (2009), ISSN 1549-9626, URL `http://dx.doi.org/10.1021/ct9004454`.

[47] O. Christiansen, J. Chem. Phys. **120**, 2149–2159 (2004), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/1.1637579`.

[48] M. J. Bramley, W. H. Green, and N. C. Handy, Mol. Phys. **73**, 1183–1208 (1991), ISSN 1362-3028, URL `http://dx.doi.org/10.1080/00268979100101871`.

[49] E. Mátyus, A. M. S. Daría, and G. Avila, Chem. Comm. **59**, 366 (2023), URL `https://doi.org/10.1039/d2cc05123k`.

[50] Z. Bačić and J. C. Light, Annu. Rev. Phys. Chem. **40**, 469 (1989).

[51] J. M. Hutson, Annu. Rev. Phys. Chem. **41**, 123 (1990).

[52] I. Simkó, C. Schran, F. Brieuc, C. Fábri, O. Asvany, S. Schlemmer, D. Marx, and A. G. Császár, Angew. Chem. Int. Ed. **62**, e202306744 (2023), URL `https://onlinelibrary.wiley.com/doi/abs/10.1002/anie.202306744`.

[53] A. Yachmenev and S. N. Yurchenko, J. Chem. Phys. **143**, 014105 (2015), URL `https://doi.org/10.1063/1.4923039`.

[54] F. Gatti, C. Iung, M. Menou, Y. Justum, A. Nauts, and X. Chapuisat, J. Chem. Phys. **108**, 8804–8820 (1998), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/1.476327`.

[55] C. Leforestier, A. Viel, F. Gatti, C. Muñoz, and C. Iung, J. Chem. Phys. **114**, 2099–2105 (2001), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/1.1337048`.

[56] M. J. Bramley and N. C. Handy, J. Chem. Phys. **98**, 1378–1397 (1993), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/1.464305`.

[57] C. Iung and F. Gatti, Int. J. Quantum Chem. **106**, 130–151 (2005), ISSN 1097-461X, URL `http://dx.doi.org/10.1002/qua.20728`.

[58] F. Gatti and C. Iung, Phys. Rep. **484**, 1–69 (2009), ISSN 0370-1573, URL `http://dx.doi.org/10.1016/j.physrep.2009.05.003`.

[59] E. L. Klinting, D. Lauvergnat, and O. Christiansen, J. Chem. Theory Comput. **16**, 4505–4520 (2020), ISSN 1549-9626, URL `http://dx.doi.org/10.1021/acs.jctc.0c00261`.

[60] J. M. Bowman, T. Carrington, and H.-D. Meyer, Mol. Phys. **106**, 2145–2182 (2008), ISSN 1362-3028, URL `http://dx.doi.org/10.1080/00268970802258609`.

[61] K. Oenen, D. F. Dinu, and K. R. Liedl, J. Chem. Phys. **160**, 014104 (2024), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/5.0180657`.

[62] D. Mendive-Tapia, H.-D. Meyer, and O. Vendrell, J. Chem. Theory Comput. **19**, 1144–1156 (2023), ISSN 1549-9626, URL `http://dx.doi.org/10.1021/acs.jctc.2c01089`.

[63] M. Schneider and G. Rauhut, J. Chem. Phys. **161**, 094102 (2024), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/5.0225991`.

[64] T. C. Thompson and D. G. Truhlar, J. Chem. Phys. **77**, 3031 (1982), URL `https://doi.org/10.1063/1.444226`.

[65] R. C. Mayrhofer and E. L. Sibert, Theor. Chim. Acta **92**, 107–122 (1995), ISSN 1432-2234, URL `http://dx.doi.org/10.1007/BF01134217`.

[66] J. Zúñiga, J. A. G. Picón, A. Bastida, and A. Requena, J. Chem. Phys. **122**, 224319 (2005), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/1.1929738`.

[67] J. Zúñiga, A. Bastida, and A. Requena, J. Chem. Phys. **115**, 139–148 (2001), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/1.1377893`.

[68] I. W. Bulik, M. J. Frisch, and P. H. Vaccaro, J. Chem. Phys. **147**, 044110 (2017), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/1.4995440`.

[69] J. M. Bowman, J. Zûñiga, and A. Wierzbicki, J. Chem. Phys. **90**, 2708–2713 (1989), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/1.455918`.

[70] K. Yagi, M. Keçeli, and S. Hirata, J. Chem. Phys. **137**, 204118 (2012), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/1.4767776`.

[71] B. Thomsen, K. Yagi, and O. Christiansen, J. Chem. Phys. **140**, 154102 (2014), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/1.4870775`.

[72] M. Chan, S. Manzhos, T. J. Carrington, and K. Yamashita, J. Chem. Theory Comput. **8**, 2053 (2012), URL `https://doi.org/10.1021/ct300248n`.

[73] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, J. Mach. Learn. Res. **22**, 2617 (2022), ISSN 1532-4435, URL `https://dl.acm.org/doi/abs/10.5555/3546258.3546315`.

[74] D. Rezende and S. Mohamed, in *Proceedings of the 32nd International Conference on Machine Learning, ICML*, edited by F. Bach and D. Blei (PMLR, 2015), vol. 37 of *Proceedings of Machine Learning Research*, pp. 1530–1538, URL `https://proceedings.mlr.press/v37/rezende15.html`.

[75] Y. Saleh and A. Iske, *Inducing Riesz and orthonormal bases in L^2 via composition operators*, preprint (2024), `2406.18613`.

[76] J. Behrmann, W. Grathwohl, R. T. Q. Chen, D. Duvenaud, and J.-H. Jacobsen, in *Proceedings of the 36th International Conference on Machine Learning*, edited by K. Chaudhuri and R. Salakhutdinov (PMLR, 2019), vol. 97 of *Proceedings of Machine Learning Research*, pp. 573–582, URL `https://proceedings.mlr.press/v97/behrmann19a.html`.

[77] G. Avila and T. Carrington, J. Chem. Phys. **139**, 134114 (2013), URL `https://doi.org/10.1063%2F1.4821348`.

[78] W. Yang and A. C. Peet, Chem. Phys. Lett. **153**, 98 (1988), URL `https://doi.org/10.1016/0009-2614(88)80139-8`.

[79] J. Toulouse and C. J. Umrigar, J. Chem. Phys. **128**, 174101 (2008), URL `https://doi.org/10.1063/1.2908237`.

[80] A. Cuzzocrea, A. Scemama, W. J. Briels, S. Moroni, and C. Filippi, J. Chem. Theory Comput. **16**, 4203 (2020), URL `https://pubs.acs.org/doi/full/10.1021/acs.jctc.0c00147`.

[81] K. L. Chubb, A. Yachmenev, J. Tennyson, and S. N. Yurchenko, J. Chem. Phys. **149**, 014101 (2018), URL `https://doi.org/10.1063/1.5031844`.

[82] S. N. Yurchenko and T. M. Mellor, J. Chem. Phys. **153**, 154106 (2020), URL `https://doi.org/10.1063/5.0019546`.

[83] A. A. A. Azzam, J. Tennyson, S. N. Yurchenko, and O. V. Naumenko, Mon. Not. R. Astron. Soc. **460**, 4063 (2016), URL `https://doi.org/10.1093/mnras/stw1133`.

[84] A. F. Al-Refaie, A. Yachmenev, J. Tennyson, and S. N. Yurchenko, Mon. Not. R. Astron. Soc. **448**, 1704 (2015).

[85] T. Van Mourik, G. J. Harris, O. L. Polyansky, J. Tennyson, A. G. Császár, and P. J. Knowles, J. Chem. Phys. **115**, 3706 (2001).

[86] P. M. Felker and Z. Bačić, J. Nanophotonics **151**, 024305 (2019), ISSN 0021-9606, URL `https://doi.org/10.1063/1.5111131`.

[87] Y. Saleh, Dissertation, Universität Hamburg, Hamburg, Germany (2023), URL `https://ediss.sub.uni-hamburg.de/handle/ediss/10390`.

[88] B. Schröder and G. Rauhut, *Vibrational Configuration Interaction Theory* (WORLD SCIENTIFIC, 2021), pp. 1–40, 0, ISBN 978-981-12-3790-4, URL `https://doi.org/10.1142/9789811237911_0001`.

[89] P. Bilous, A. Pálffy, and F. Marquardt, Phys. Rev. Lett. **131**, 133002 (2023), URL `https://link.aps.org/doi/10.1103/PhysRevLett.131.133002`.

[90] T. Petrenko and G. Rauhut, J. Chem. Phys. **146** (2017).

[91] Y. Saleh, Á. F. Corral, E. Vogt, A. Iske, J. Küpper, and A. Yachmenev, *Computing excited states of molecules using normalizing flows*, preprint (2023), `2308.16468`.

[92] Q. Zhang, R.-S. Wang, and L. Wang, J. Chem. Phys. **161**, 024103 (2024), URL `10.1063/5.0209255`.

[93] M. Entwistle, Z. Schätzle, P. A. Erdman, J. Hermann, and F. Noé, Nat. Commun. **14**, 274 (2023), `2203.09472`.

[94] D. Gottlieb and S. A. Orszag, *Numerical analysis of spectral methods: theory and applications* (SIAM, 1977), URL `https://epubs.siam.org/doi/book/10.1137/1.9781611970425`.

[95] K. Cranmer, S. Golkar, and D. Pappadopulo, arXiv preprint arXiv:1904.05903 (2019), `1904.05903`.

[96] Y. Saleh, A. Iske, A. Yachmenev, and J. Küpper, Proc. Appl. Math. Mech. **23**, e202200239 (2023), `2212.01383`, URL `https://onlinelibrary.wiley.com/doi/full/10.1002/pamm.202200239`.

[97] R. Fisher, *Statistical Methods for Research Workers*, Biological monographs and manuals (Oliver and Boyd, 1925), URL `https://books.google.de/books?id=IONBAAAAIAAJ`.

[98] T. H. Dunning, J. Chem. Phys. **90**, 1007 (1989), URL `https://dx.doi.org/10.1063/1.456153`.

[99] R. A. Kendall, T. H. Dunning, Jr., and R. J. Harrison, J. Chem. Phys. **96**, 6796 (1992), URL `http://dx.doi.org/10.1063/1.462569`.

[100] D. G. A. Smith, L. A. Burns, A. C. Simmonett, R. M. Parrish, M. C. Schieber, R. Galvelis, P. Kraus, H. Kruse, R. D. Remigio, A. Alenaizan, et al., J. Chem. Phys. **152**, 184108 (2020), URL `https://doi.org/10.1063/5.0006002`.

[101] V. Lebedev and D. Laikov, Dokl. Math. **59**, 477 (1999).

[102] A. D. Becke, J. Chem. Phys. **88**, 2547–2553 (1988), ISSN 1089-7690, URL `http://dx.doi.org/10.1063/1.454033`.

[103] F. M. Fernández and J. Garcia, ChemistrySelect **6**, 9527–9534 (2021), ISSN 2365-6549, URL `http://dx.doi.org/10.1002/slct.202102509`.

[104] B. J. Braams and J. M. Bowman, Int. Rev. Phys. Chem. **28**, 577 (2009), URL `https://doi.org/10.1080/01442350903234923`.

[105] S. Manzhos and T. Carrington, Chem. Rev. **121**, 10187 (2021), URL `https://doi.org/10.1021/acs.chemrev.0c00665`.

[106] I. I. Mizus, A. A. Kyuberis, N. F. Zobov, V. Y. Makhnev, O. L. Polyansky, and J. Tennyson, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **376**, 20170149 (2018), URL `https://doi.org/10.1098/rsta.2017.0149`.

[107] Y. Saleh, V. Sanjay, A. Iske, A. Yachmenev, and J. Küpper, J. Chem. Phys. **155**, 144109 (2021), `2104.00708`, URL `https://doi.org/10.1063/5.0057051`.

[108] P. R. Franke, J. F. Stanton, and G. E. Douberly, J. Phys. Chem. A **125**, 1301 (2021), URL `https://doi.org/10.1021/acs.jpca.0c09526`.

[109] M. Mendolicchio, J. Bloino, and V. Barone, J. Chem. Theory Comput. **18**, 7603 (2022), URL `https://doi.org/10.1021/acs.jctc.2c00773`.

[110] X. Chapuisat, A. Nauts, and J.-P. Brunet, Mol. Phys. **72**, 1 (1991), URL `https://doi.org/10.1080/00268979100100011`.

[111] B. Ziegler and G. Rauhut, J. Chem. Theory Comput. **15**, 4187–4196 (2019), ISSN 1549-9626, URL `http://dx.doi.org/10.1021/acs.jctc.9b00381`.

[112] X.-G. Wang and T. Carrington, J. Chem. Phys. **130**, 094101 (2009), URL `https://doi.org/10.1063%2F1.3077130`.

[113] Y. Saleh, Á. Fernández Corral, E. Vogt, A. Iske, J. Küpper, and A. Yachmenev, J. Chem. Theory Comput. **21**, 5221 (2025), URL `https://doi.org/10.1021/acs.jctc.5c00590`.

[114] Q. Zhang, X. Wang, R. Shi, X. Ren, H. Wang, and L. Wang, *Neural canonical transformations for quantum anharmonic solids of lithium* (2024), `2412.12451`, URL `https://arxiv.org/abs/2412.12451`.

[115] DeepMind, I. Babuschkin, K. Baumli, A. Bell, S. Bhupatiraju, J. Bruce, P. Buchlovsky, D. Budden, T. Cai, A. Clark, et al., *The DeepMind JAX Ecosystem* (2020), URL `http://github.com/deepmind`.

[116] E. K. Conway, I. E. Gordon, J. Tennyson, O. L. Polyansky, S. N. Yurchenko, and K. Chance, Atmos. Chem. Phys., **20**, 10015 (2020), URL `https://acp.copernicus.org/articles/20/10015/2020/`.

[117] P. M. Morse, Phys. Rev. **34**, 57 (1929), URL `https://link.aps.org/doi/10.1103/PhysRev.34.57`.

[118] E. Vogt, D. S. Sage, and H. G. Kjaergaard, Mol. Phys. **117**, 1629 (2019), URL `https://doi.org/10.1080/00268976.2018.1521529`.

[119] S. N. Yurchenko, A. Yachmenev, and R. I. Ovsyannikov, J. Chem. Theory Comput. **13**, 4368 (2017), `1708.07185`, URL `https://doi.org/10.1021/acs.jctc.7b00506`.

[120] J. P. Boyd, *Chebyshev and Fourier Spectral Methods*, Applied Mathematical Sciences (Springer, 2000).

[121] E. Vogt, Á. F. Corral, Y. Saleh, and A. Yachmenev (2025), `2502.15750`.

[122] H. Xie, L. Zhang, and L. Wang (2021), `2105.08644`.

[123] I. Kobyzev, S. J. Prince, and M. A. Brubaker, IEEE Trans. Pattern Anal. Mach. Intell. **43**, 3964 (2020).

[124] T. M. Apostol, *Mathematical analysis; 2nd ed.*, Addison-Wesley series in mathematics (Addison-Wesley, 1974).

[125] W. Rudin, *Principles of Mathematical Analysis*, International series in pure and applied mathematics (McGraw-Hill, 1976), ISBN 9780070856134, URL `https://books.google.de/books?id=kwqzPAAACAAJ`.

[126] P. Jaming and I. Simon, Bulletin des Sciences Mathématiques **166**, 102933 (2021), ISSN 0007-4497, URL `https://www.sciencedirect.com/science/article/pii/S0007449720301032`.

[127] P. Jensen, J. Mol. Spectrosc. **128**, 478 (1988), URL `https://www.sciencedirect.com/science/article/pii/0022285288901646`.

[128] A. Yachmenev, S. N. Yurchenko, P. Jensen, and W. Thiel, J. Chem. Phys. **134**, 244307 (2011), URL `https://doi.org/10.1063/1.3599927`.

[129] P. Jensen, J. Mol. Spectrosc. **133**, 438 (1989), URL `https://www.sciencedirect.com/science/article/pii/0022285289902038`.

[130] V. G. Tyuterev, S. A. Tashkun, and D. W. Schwenke, Chem. Phys. Lett. **348**, 223 (2001), URL `https://www.sciencedirect.com/science/article/pii/S0009261401010934`.

[131] T. Morawietz, V. Sharma, and J. Behler, J. Chem. Phys. **136**, 064103 (2012), URL `https://aip.scitation.org/doi/full/10.1063/1.3682557`.

[132] S. K. Natarajan, T. Morawietz, and J. Behler, Phys. Chem. Chem. Phys. **17**, 8356 (2015), URL `https://pubs.rsc.org/ko/content/articlehtml/2015/cp/c4cp04751f`.

[133] S. Manzhos, R. Dawes, and T. Carrington, Int. J. Quantum Chem. **115**, 1012 (2014), URL `https://doi.org/10.1002/qua.24795`.

[134] C. I. Blaga, J. Xu, A. D. DiChiara, E. Sistrunk, K. Zhang, P. Agostini, T. A. Miller, L. F. DiMauro, and C. D. Lin, Nature **483**, 194 (2012), URL `https://www.nature.com/articles/nature10820`.

[135] M. G. Pullen, B. Wolter, A.-T. Le, M. Baudisch, M. Hemmer, A. Senftleben, C. D. Schroter, J. Ullrich, R. Moshammer, C. D. Lin, et al., Nat. Commun. **6**, 7262 (2015), `1503.03294`, URL `http://www.nature.com/ncomms/2015/150624/ncomms8262/full/ncomms8262.html`.

[136] U. D. Giovannini, J. Küpper, and A. Trabattoni, J. Phys. B **56**, 054002 (2023), URL `https://dx.doi.org/10.1088/1361-6455/acb872`.

[137] E. T. Karamatskos, G. Goldsztejn, S. Raabe, P. Stammer, T. Mullins, A. Trabattoni, R. R. Johansen, H. Stapelfeldt, S. Trippel, M. J. J. Vrakking, et al., J. Chem. Phys. **150**, 244301 (2019), `1905.03541`, URL `https://doi.org/10.1063/1.5093959`.

[138] J. Wiese, J. Onvlee, S. Trippel, and J. Küpper, Phys. Rev. Research **3**, 013089 (2020), `2003.02116`, URL `https://doi.org/10.1103/PhysRevResearch.3.013089`.

[139] M. Meckel, D. Comtois, D. Zeidler, A. Staudte, D. Pavičić, H. C. Bandulet, H. Pépin, J. C. Kieffer, R. Dörner, D. M. Villeneuve, et al., Science **320**, 1478 (2008), ISSN 0036-8075, URL `http://science.sciencemag.org/content/320/5882/1478`.

[140] M. Lewenstein, P. Balcou, M. Y. Ivanov, A. L'Huillier, and P. B. Corkum, Phys. Rev. A **49**, 2117 (1994), URL `https://journals.aps.org/pra/abstract/10.1103/PhysRevA.49.2117`.

[141] P. M. Paul, E. S. Toma, P. Breger, G. Mullot, F. Augé, P. Balcou, H. G. Muller, and P. Agostini, Science **292**, 1689 (2001), ISSN 0036-8075, URL `http://science.sciencemag.org/content/292/5522/1689`.

[142] A. L'Huillier, D. Descamps, A. Johansson, J. Norin, J. Mauritsson, and C. G. Wahlstrom, Eur. Phys. J. D **26**, 91 (2003), URL `http://www.springerlink.com/openurl.asp?genre=article&id=doi:10.1140/epjd/e2003-00072-2`.

[143] T. Zuo, A. D. Bandrauk, and P. B. Corkum, Chem. Phys. Lett. **259**, 313 (1996).

[144] Y. Ito, C. Wang, A.-T. Le, M. Okunishi, D. Ding, C. D. Lin, and K. Ueda, Struct. Dyn. **3**, 034303 (2016), URL `https://doi.org/10.1063/1.4952602`.

[145] B. Walker, B. Sheehy, K. C. Kulander, and L. F. DiMauro, Phys. Rev. Lett. **77**, 5031 (1996), URL `https://link.aps.org/doi/10.1103/PhysRevLett.77.5031`.

[146] Z. Chen, A.-T. Le, T. Morishita, and C. D. Lin, Phys. Rev. A **79**, 033409 (2009), URL `https://link.aps.org/doi/10.1103/PhysRevA.79.033409`.

[147] C. D. Lin, A.-T. Le, Z. Chen, T. Morishita, and R. Lucchese, J. Phys. B **43**, 122001 (2010), URL `http://stacks.iop.org/0953-4075/43/i=12/a=122001`.

[148] J. Xu, Z. Chen, A.-T. Le, and C. D. Lin, Phys. Rev. A **82**, 033403 (2010), URL `https://link.aps.org/doi/10.1103/PhysRevA.82.033403`.

[149] L. Schåfer, Applied Spectroscopy **30**, 123 (1976), URL `https://doi.org/10.1366/000370276774456381`.

[150] P. D. McCaffrey, J. K. Dewhurst, D. W. H. Rankin, R. J. Mawhorter, and S. Sharma, J. Chem. Phys. **128**, 204304 (2008), ISSN 0021-9606, `https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.2928806/13482931/204304_1_online.pdf`, URL `https://doi.org/10.1063/1.2928806`.

[151] M. Y. Ivanov, M. Spanner, and O. Smirnova, J. Mod. Opt. **52**, 165 (2005), URL `https://doi.org/10.1080/0950034042000275360`.

[152] N. I. Shvetsov-Shilovski, M. Lein, L. B. Madsen, E. Räsänen, C. Lemell, J. Burgdörfer, D. G. Arbó, and K. Tőkési, Phys. Rev. A **94**, 013415 (2016), URL `https://doi.org/10.1103/PhysRevA.94.013415`.

[153] P. B. Corkum, Phys. Rev. Lett. **71**, 1994 (1993), URL `https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.71.1994`.

[154] K. C. Kulander, K. J. Schafer, and J. L. Krause, in *Super-Intense Laser-Atom Physics*, edited by A. L'Huillier, B. Piraux, and K. Rzazewski (Plenum Press, New York, 1993), vol. 316 of *Nato Science Series B: Physics*, pp. 95–110, ISBN 978-1-4615-7965-6.

[155] D. M. Wolkow, Zeitschrift für Physik **94** (1935), URL `https://doi.org/10.1007/BF01331022`.

[156] L. S. Wang, J. E. Reutt, Y. T. Lee, and D. A. Shirley, J. Electron. Spectrosc. Relat. Phenom. **47**, 167 (1988), URL `http://www.sciencedirect.com/science/article/pii/0368204888850102`.

[157] L. Holmegaard, J. L. Hansen, L. Kalhøj, S. L. Kragh, H. Stapelfeldt, F. Filsinger, J. Küpper, G. Meijer, D. Dimitrovski, M. Abu-samha, et al., Nat. Phys. **6**, 428 (2010), `1003.4634`, URL `https://www.nature.com/articles/nphys1666`.

[158] A. Kværnø, BIT Numer. Math. **44**, 489 (2004).

[159] M. Okunishi, T. Morishita, G. Prümper, K. Shimada, C. D. Lin, S. Watanabe, and K. Ueda, Phys. Rev. Lett. **100**, 143001 (2008), URL `https://link.aps.org/doi/10.1103/PhysRevLett.100.143001`.

[160] A. Trabattoni, J. Wiese, U. De Giovannini, J.-F. Olivieri, T. Mullins, J. Onvlee, S.-K. Son, B. Frusteri, A. Rubio, S. Trippel, et al., Nat. Commun. **11**, 2546 (2020), `1802.06622`, URL `https://doi.org/10.1038/s41467-020-16270-0`.

[161] R. Bonham and M. Fink, *High Energy Electron Scattering* (Van Nostrand Reinhold Company, 1974), ISBN 0-442-30891-4.

[162] F. Salvat, A. Jablonski, and C. Powell, Comp. Phys. Comm. **165**, 157 (2005), ISSN 0010-4655, URL `http://www.sciencedirect.com/science/article/pii/S0010465504004795`.

[163] S.-F. Zhao, J. Xu, C. Jin, A.-T. Le, and C. D. Lin, J. Phys. B: At. Mol. Opt. Phys. **44**, 035601 (2011), ISSN 0953-4075.

[164] X. M. Tong and C. D. Lin, J. Phys. B **38**, 2593 (2005), URL `https://doi.org/10.1088/0953-4075/38/15/001`.

[165] R. Murray, M. Spanner, S. Patchkovskii, and M. Y. Ivanov, Phys. Rev. Lett. **106**, 173001 (2011), URL `https://link.aps.org/doi/10.1103/PhysRevLett.106.173001`.

[166] M. Li, J.-W. Geng, M. Han, M. M. Liu, L.-Y. Peng, Q. Gong, and Y. Liu, Phys. Rev. A **93**, 013402 (2016), URL `https://doi.org/10.1103/PhysRevA.93.013402`.

[167] G. Cybenko, Math. Control Signals Syst. **2**, 303 (1989).

[168] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, et al., *JAX: composable transformations of Python+NumPy programs* (2018), URL `http://github.com/jax-ml/jax`.

[169] P. Kidger, Dissertation, University of Oxford (2021), URL `https://arxiv.org/abs/2202.02435`.

[170] A. T. J. B. Eppink and D. H. Parker, Rev. Sci. Instrum. **68**, 3477 (1997), URL `https://aip.scitation.org/doi/abs/10.1063/1.1148310`.

[171] A. I. Chichinin, K. H. Gericke, S. Kauczok, and C. Maul, Int. Rev. Phys. Chem. **28**, 607 (2009), URL `http://www.tandfonline.com/doi/abs/10.1080/01442350903235045`.

[172] S. Cuomo, V. S. Di Cola, F. Giampaolo, G. Rozza, M. Raissi, and F. Piccialli, J. Sci. Comput. **92**, 88 (2022), URL `https://doi.org/10.1007/s10915-022-01939-z`.

[173] X. Jin, S. Cai, H. Li, and G. E. Karniadakis, J. Comput. Phys. **426**, 109951 (2021), URL `https://doi.org/10.1016%2Fj.jcp.2020.109951`.

[174] H. Eivazi, M. Tahani, P. Schlatter, and R. Vinuesa, Phys. Fluids **34**, 075117 (2022), ISSN 1070-6631, URL `https://doi.org/10.1063/5.0095270`.

[175] C. Rao, H. Sun, and Y. Liu, Theor. Appl. Mech. Lett. **10**, 207 (2020).

[176] K. Shah, P. Stiller, N. Hoffmann, and A. Cangi, *Physics-informed neural networks as solvers for the time-dependent Schrödinger equation* (2022), `2210.12522`.

[177] Z. Chen, Y. Liu, and H. Sun, Nat. Commun. **12** (2021), ISSN 2041-1723, URL `http://dx.doi.org/10.1038/s41467-021-26434-1`.

[178] H. Xu, D. Zhang, and N. Wang, J. Comput. Phys. **445**, 110592 (2021), ISSN 0021-9991, URL `http://dx.doi.org/10.1016/j.jcp.2021.110592`.

[179] H. Xu and D. Zhang, Phys. Rev. Research **3**, 033270 (2021), ISSN 2643-1564, URL `http://dx.doi.org/10.1103/PhysRevResearch.3.033270`.

[180] P. Sharma, L. Evans, M. Tindall, and P. Nithiarasu, Arch. Comput. Methods Eng. **30**, 2929 (2023).

[181] E. Hairer and G. Wanner, *Solving Ordinary Differential Equations II. Stiff and Differential-Algebraic Problems*, vol. 14 (Springer Berlin, Heidelberg, 1996).

[182] S. Kim, W. Ji, S. Deng, Y. Ma, and C. Rackauckas, Chaos **31**, 093122 (2021), ISSN 1089-7682, URL `http://dx.doi.org/10.1063/5.0060697`.

[183] C. Moya and G. Lin, Neural. Comput. Appl. **35**, 3789 (2022), URL `https://doi.org/10.1007/s00521-022-07886-y`.

[184] M. Antoñana, E. Alberdi, J. Makazaga, and A. Murua, Celestial Mech. & Dyn. Astron. **134** (2022), URL `http://dx.doi.org/10.1007/s10569-022-10081-9`.

[185] H. Goldstein, *Classical mechanics* (Pearson Education, Amsterdam, 2002), ISBN 978-8-177-58283-3.

[186] N. I. Shvetsov-Shilovski, Eur. Phys. J. D **75**, 130 (2021), URL `http://dx.doi.org/10.1140/epjd/s10053-021-00134-3`.

[187] J. Wiese, Dissertation, Universität Hamburg, Hamburg, Germany (2020), URL `https://ediss.sub.uni-hamburg.de/handle/ediss/8523`.

[188] J. Toulouse, A. Savin, and H.-J. Flad, Int. J. Quantum Chem. **100**, 1047 (2004), `https://onlinelibrary.wiley.com/doi/pdf/10.1002/qua.20259`, URL `https://onlinelibrary.wiley.com/doi/abs/10.1002/qua.20259`.

[189] A. Iserles, *A First Course in the Numerical Analysis of Differential Equations* (Cambridge University Press, 2008), 2nd ed., ISBN 9780511995569, URL `http://dx.doi.org/10.1017/CBO9780511995569`.

[190] S. González Pinto, S. Pérez Rodríguez, and J. Montijano Torcal, J. Comput. & Appl. Math. **82**, 129 (1997), ISSN 0377-0427, 7th ICCAM 96 Congress, URL `https://www.sciencedirect.com/science/article/pii/S0377042797000861`.

[191] E. Hairer and G. Wanner, J. Comput. & Appl. Math. **111**, 93 (1999), ISSN 0377-0427, URL `https://www.sciencedirect.com/science/article/pii/S037704279900134X`.

[192] C. A. Kennedy and M. H. Carpenter, Appl. Num. Math. **146**, 221 (2019), ISSN 0168-9274, URL `https://www.sciencedirect.com/science/article/pii/S0168927419301801`.

[193] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, et al., *JAX: composable transformations of Python+NumPy programs* (2018), URL `http://github.com/google/jax`.

[194] J. Heek, A. Levskaya, A. Oliver, M. Ritter, B. Rondepierre, A. Steiner, and M. van Zee, *Flax: A neural network library and ecosystem for JAX* (2023), URL `http://github.com/google/flax`.

[195] D. P. Kingma and J. Ba, arXiv preprint arXiv:1412.6980 (2015), URL `http://arxiv.org/abs/1412.6980`.

[196] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, SIAM J. Sci. Comp. **16**, 1190 (1995), URL `https://doi.org/10.1137/0916069`.

[197] M. Blondel, Q. Berthet, M. Cuturi, R. Frostig, S. Hoyer, F. Llinares-López, F. Pedregosa, and J.-P. Vert, in *Advances in Neural Information Processing Systems (NIPS)* (Curran Associates, Inc., 2022), vol. 35, pp. 5230–5242, `2105.15183`, URL `https://proceedings.neurips.cc/paper_files/paper/2022/file/228b9279ecf9bbafe582406850c57115-Paper-Conference.pdf`.

[198] M. Panicker and C. Babu, IOSR J. Eng. **02**, 1352 (2012).

[199] R. C. Daileda, *The two dimensional heat equation* (2012), URL `http://ramanujan.math.trinity.edu/rdaileda/teach/s12/m3357/lectures/lecture_3_6_short.pdf`.

[200] G. I. Taylor and A. E. Green, Proc. Royal Soc. London A **158**, 499 (1937), URL `https://doi.org/10.1098/rspa.1937.0036`.

[201] J. Gallier, *Remarks on the cayley representation of orthogonal matrices and on perturbing the diagonal of a matrix to make it invertible* (2013), `math/0606320`, URL `https://arxiv.org/abs/math/0606320`.

[202] P. A. M. Dirac, Mathematical Proceedings of the Cambridge Philosophical Society **26**, 376–385 (1930).

[203] J. Frenkel, *Wave Mechanics: Advanced General Theory*, International series of monographs on physics (Clarendon Press, 1934), URL `https://books.google.de/books?id=ZGcmAAAAMAAJ`.

[204] A. McLachlan, Molecular Physics **8**, 39 (1964), URL `https://doi.org/10.1080/00268976400100041`.