White Paper: “Advancing Quantum Mechanics with Mathematics and Statistics” (IPAM Long Program, Spring 2022)

Mihail Bogojeski, Eric Cances, Stefan Chmiela, Brian Ernst, Fabian M. Faulstich, Carlos García-Cervera, Szabolcs Góger, Matteo Gori, Muhammad Hassan, Scott Jensen, Almaz Khabibrakhmanov, Philip Kurian, Rafael Lainez Reyes, Marcel Langer, Remi Leano, Kangbo Li, Lin Lin, Vincent Martinetto, Abigail Poteshman, Trine Kay Quady, Zeno Schätzle, Zachary Sparrow, Martin Stöhr, Kevin D Stubbs, Lucas Tecot, Alexandre Tkatchenko, and Guido Falk von Rudorff

Table of Contents:

1. Executive Summary

This document summarizes the activities and outcomes of the Long Program “Advancing Quantum Mechanics with Mathematics and Statistics” which was held at the Institute of Pure and Applied Mathematics (IPAM) from March 7 to June 10, 2022. It also briefly explores some of the current open questions and future directions in the field of electronic structure theory and computational chemistry as well as related fields that were discussed during the program.

Field Theory Approaches: One challenge in quantum mechanics for large systems is connecting established methods across scales. Quantum Field Theory (QFT) provides a natural framework to derive existing methods from first principles, highlighting common approximations and a pathway to more accurate ones. This framework allows the description of collective and long-range effects in large complexes. In this context, the development of a rigorous effective field theory approach and the analysis of the light-matter interaction beyond the semiclassical approximations in large molecular systems have been identified as key open questions.

Excited and open quantum systems: Multiscale approaches to quantum mechanics need to describe collective degrees of freedom arising from large time and length scales. Open quantum system theory offers a productive way to describe excited electronic states, especially in the context of light-matter interactions. This accurate description is a key aspect both for cavity molecular quantum electrodynamics and large biomolecular complexes. The development of rigorous, accurate, and efficient methodologies for multiscale modeling across optical, electronic, and vibrational degrees of freedom has been envisaged as one of the main directions for future research.

Embedding theories: Embedding theories provide an efficient way to combine computational methods of different scales to treat large quantum systems. Within this program, a variety of embedding methods were discussed and scrutinized. Particular attention has been given to density matrix embedding theory, for which a numerical remedy for the chemical potential fitting problem has been developed. The implementation thereof is important as it extends the numerical applicability of embedding methods to larger and more complicated systems.

Response functions in molecules and solids: The number of materials and molecules grows combinatorially with the elements and building blocks considered. Dominating quantum effects can often be recovered with response functions describing how electronic structure changes under external influences. Participants devised new methods to obtain some of these response functions relevant to interacting systems and many-body interactions. Open challenges include obtaining high-quality response functions and the systematic improvement over current limitations of quantum mechanical calculations.

Non-covalent Interactions: Macroscopic properties of matter are chiefly determined by intermolecular interactions. In large systems, non-covalent interactions are highly non-local with van der Waals (vdW) dispersion being of particular importance. Due to its long-range nature and complex quantum many-body character, an accurate treatment of the vdW dispersion is currently challenging for large systems. During the program an interdisciplinary effort addressed the current lack of accurate and efficient methods capturing vdW dispersion interactions. In addition to the development of novel numerical approaches, this multidisciplinary perspective allows for a better understanding of the physical origins of non-covalent interactions.

Machine Learning for Quantum Mechanics: Quantum mechanical problems are high-dimensional and often nonlinear. If machine learning methods can yield models of relevant accuracy, more applications and systems become feasible. The program explored multiple complementary methods that aim to describe long-range interactions and improve transferability by introducing physical constraints directly. The participants identified interpretability of models, inclusion of functional derivatives, and treatment of large-scale many-body systems as some of the key open questions.

Additional Long Program Activity - 2D Moire Materials Workshop: Since its isolation by Geim, Novoselov, and collaborators in 2004, single layer graphene has drawn intense interest for its remarkable physical and theoretical properties. Graphene has the highest known tensile strength, and conductivity. More recently, theoretical studies have focused on the mathematical modeling of twisted bilayer graphene (TBG), a metamaterial constructed by stacking two layers of graphene and then rotating one layer with respect to the other. This has led to the development of models such as the Bistritzer-Mcdonald model (2011) and the Chiral model proposed by Tarnopolsky and coworkers (2019). The working group on 2D materials focused on reading and understanding the most recent mathematical results on the spectral properties of TBG and strained graphene at magic angles.

2. Field theory approaches

Introduction: Quantum field theory (QFT) is a natural language to describe how quantum effects manifest on a wide range of length scales, from atomic to supra- and intermolecular. It allows us to tackle and overcome some significant issues that arise in scaling up existing methods for solving electronic structure problems. In the usual quantum mechanical description of a particle system, the positions and momenta of each electron are quantized, i.e. they become operators acting on state vectors in N-particle Hilbert space. In the electronic structure problem, this results in a description of the quantum state of the whole system based on a many-body antisymmetric wavefunction, whose domain is a very high-dimensional space with few similarities to three-dimensional position coordinate space. Furthermore, QFT provides an approach to quantum many-body systems that is analogous to the Eulerian description of fluid dynamics: the relevant (quantized) degrees of freedom are not the positions and momenta of the single electrons, but the amplitudes of the Schrödinger field at each point in real space. This allows us to study correlations and structure formations in position space because the spatial coordinates play the role of a continuum of labels (rather than discrete eigenvalues of the position operators), and to describe the dynamics of the field operators in the Heisenberg representation at each point in space, with an arbitrary level of resolution.

Non-relativistic QFT provides a means to derive DFT from first principles through establishment of a semiclassical approximation method as the effective action (EA). Despite what such a QFT-based approach to DFT could potentially offer in terms of conceptual insight and further development of the DFT framework, it has remained quite underused in the quantum chemistry community. A part of the working group's activity has been devoted to rediscovering this approach and envisaging further developments and refinement of the DFT method. Studying quantum mechanics of electronic structure in million-atom systems would require providing an accurate and efficient description of the interaction of matter with the electromagnetic field. The contribution due to radiative electromagnetic interactions among electrons becomes non-negligible for large molecular complexes, whose length scale is beyond the order of 100 nm. In this regime, a self-consistent description of the dynamics of the coupled system of matter and the electromagnetic field is required. The long-range, non-Debye-screened electrodynamical interactions are central in the emergence of collective phenomena in macroscopic quantum systems. This aspect is fundamental in the perspective of cavity quantum electrodynamics (QED) for molecular systems—highlighted by speakers Angel Rubio and

Prineha Narang in Workshop I—that represents a promising field both for testing the prediction of the theory and for developing new methods to control chemical properties in molecular systems. Finally, QFT methods can be crucial to develop a promising and accurate description of long-range correlation energy formulated in terms of the adiabatic connection fluctuation-dissipation theorem applied to the polarizability density field (see section 5 and section 6). An effective field theory description of the polarization field derived from the first principles of molecular QED would provide a major contribution to developing efficient numerical implementations for this approach.

Field theory formulation of DFT: DFT represents a diverse array of workhorse numerical approaches for solving electronic structure problems, especially in large molecular systems. Part of the success of DFT, as compared to other methods iis its reliance on a classical density field over real space to describe quantum electronic structure from atomic to molecular scales. DFT in its original formulation provides an exact solution to the many-body electron problem in quantum chemistry. Practically, DFT's application requires the approximation of the exchange-correlation functional. Much attention has been devoted to the analysis of different approximations to assess their mutual consistency. Even generalizations of DFT have been proposed to overcome the problem of the exchange-correlation functional, as discussed by Kieron Burke in his talk during Workshop IV. However, a rigorous derivation of DFT — employing the widespread local density approximation (LDA) — from non-relativistic QED first principles has received less attention from the greater quantum chemistry community. In fact, the few existing articles tackling this problem have a relatively low number of citations (no more than 60 in total) despite being published more than two to four decades ago. In these works, an effective field theory approach treating the "electron gas" as a fluid has been used to derive the exact DFT energy functional, in which it is possible to control the level of the approximation and to move beyond the Kohn-Sham scheme and LDA. The IPAM Long Program provided the opportune moment to rediscover these 'sleeping beauty' papers, from which new research lines have been envisioned to implement effective action approaches for DFT in an orbital-free framework. These research lines will complement other promising techniques such as auxiliary-field quantum Monte Carlo, which were presented by Shiwei Zhang and Sandeep Sharma in Workshop IV.

Light-matter interaction: The interaction of light and matter is at the foundation of a number of exciting research areas in chemistry and physics. A coupled system of light and matter, often through the use of a cavity, can be significantly distinct from an isolated system of matter in vacuum. For quantum chemistry, coupling photons and molecules can change the molecular energy landscape, thus potentially modifying chemical reactions. For materials, light can alter the ground state of matter, potentially inducing phases that are otherwise not manifest in the vacuum ground state. Such phases include superconductivity, ferroelectricity, and the anomalous quantum Hall effect. If the light can be precisely controlled through lasers or microwaves, the interaction of light and matter can be used to precisely control the dynamics of the matter. This is at the heart of qubit control, which is an essential part of quantum computing. These and related applications have strongly motivated members of the Field Theory Approaches working group to understand the quantum theory of light and matter.

Work done at IPAM: During the IPAM Long Program, the group discussed the basic components that are used to build up the quantum theory of light and matter. We started with the classical Lagrangian theory of fields, from which the quantum theory of fields is formulated through the canonical quantization of the field Lagrangian. We then proceeded to understand the classical and semiclassical theory of the interaction between light and matter. Here we studied the minimum coupling Hamiltonian and derived the length gauge representation by introducing the dipole approximation and a gauge transformation. Finally, both matter and electromagnetic field are quantized, which led us to the commonly used working model for cavity QED.

Outlook: A number of possible research directions were discussed by the group. Incorporating QFT approaches into DFT at the formulational level might provide new ways of approximating exchange-correlation potentials. Despite some selected works regarding the study of quantum systems with quantized fields, much remains to be done, motivated by the recent interest and new applications within the physics and materials science communities. The treatment of quantized fields introduces additional challenges to the development of efficient numerical methodologies for the study of electronic structures within the context of QED. In addition, optimal control of matter with light beyond the Jaynes-Cummings model could have significant impact in different areas, such as the development of quantum computers.

3. Excited and open quantum systems

Introduction: The Excited and Open Quantum Systems working group was guided by the overarching objective of bridging time and length scales for complex systems involving many different degrees of freedom. Beginning with the commonly used models of two-level systems in quantum optics, we then proceeded to study ensemble density functional theory and Markovian and non-Markovian approaches to open quantum systems, culminating in research topics centered on the integration of excited states from high-level quantum chemistry calculations into classical molecular dynamics packages.

Two-level systems: Discrete two-level systems (TLS) have been widely used to describe the interaction of quantized matter with the electromagnetic field. Within the setting of restricted Hilbert spaces, where atoms to complex molecules can be treated effectively as such TLS, matter absorbs excitations from, or emits excitations to, particular quantized mode(s) of the field. In the canonical Jaynes-Cummings (JC) model of quantum optics, the electromagnetic field is restricted to a single mode, and the JC solution (exhibiting so-called Rabi oscillations) is presented for a single TLS in the rotating wave approximation, where the fast-oscillating ("counter-rotating") terms are filtered out. The JC model can be extended to the Tavis-Cummings (TC) model, originally proposed by Robert Dicke, where many TLS interact with a single mode of the field. During Workshop 1 of the QMM long program, several speakers—including Angel Rubio ("Polaritonic quantum materials: A first-principles QEDFT perspective"), Prineha Narang, and Hakan Tureci ("Quantum electrodynamic modeling of superconducting quantum systems")—referred to the TC model or augmented it by including either (1) more levels to each quantized matter system; (2) more modes of the electromagnetic field; (3) more multi-level matter systems; (4) distinct modes of a different field (such as for phonons); or (5) combinations of these. By the end of the long program, the core participants in the Excited and Open Quantum Systems working group reached the point of engaging research-level questions in the domain of so-called "cutoff-free" quantum electrodynamics (QED), where all modes of the quantized electromagnetic field are included without the specter of infinite ultraviolet divergences.

Ensemble DFT: The original papers establishing the foundations of Ensemble DFT (EDFT) were published by Gross, Oliveira, and Kohn in 1988. The basis of ensemble DFT is a theorem by Theophilou for quantum mechanics and an excited-state variational principle. The Excited and Open Quantum Systems working group discussed the derivations contained within the first two of the three original papers published by Gross, Oliveira, and Kohn, which established EDFT as a method. The third of these papers detailed the calculation for the helium (He) atom, which yielded unsatisfactory results largely overestimating the excitation energies of He. Further progress within the field includes the identification of one method for correcting 'ghost' interactions in the Hartree and exchange terms by Gidopolus and Gross, which greatly improved the accuracy of calculations. In 2014, the work of Zeng-hui Yang, Aurora Pribram-Jones, and others demonstrated how a different choice of ensemble (the GOK II ensemble), which provided a simpler analytic form for excitation energies, allowed for the inclusion of double excitations, which are missed in practice by current time-dependent DFT calculations. In recent years, the work of other theory developers (Fromager, Pittalis, Gould et al.) on the construction of improved ensemble correlation functionals have enabled a deeper fundamental understanding of EDFT. Currently, steps are being taken to implement EDFT into major codes, which will enable the study of additional phenomena, such as conical intersections.

Open Quantum Systems: The study of large systems is very challenging. However, we are often only interested in the dynamics of a subsystem of a larger system. The main question in such a case is how one can obtain an effective description of the embedded subsystem, in a way that captures the interaction with the larger system, which is treated as a reservoir. Two examples of such formalisms for reduction of degrees of freedom are those introduced by Mori and Zwanzig, and by Lindblad. The Mori-Zwanzig approach leads to an exact, non-Markovian, generalized Langevin equation for the subsystem in which the interaction with the bath is characterized by a memory kernel. The Lindblad equation, on the other hand, is a linear, Markovian approximation to the Liouville-von Neumann (LvN) equation that describes the dynamics of the density matrix of an embedded quantum system, and is the quantum analogue of the Fokker-Planck equation. Following the work by Daniel Manzano and the book by Petruccione and Breuer, the group discussed the derivation of the Lindblad equation from first principles, starting from the LvN equation, and paying special attention to the specific approximations that are done, with the aim of understanding its applicability and limitations, as well as possible extensions.

Excited states in classical molecular dynamics (MD): The Excited and Open Quantum Systems working group concluded its discussions with a research talk on incorporating (photo)excited states from high-level quantum chemistry calculations into classical molecular dynamics simulations for discerning differences between ground and excited-state protein configurations over long times (500 ns). These MD simulations have been motivated by recent experiments showcasing the emergence of a pronounced low-frequency terahertz (THz) phonon mode in an aqueous solution of fluorophore-decorated bovine serum albumin (BSA) proteins. Such few-THz signatures in protein solutions emerge only when the sample is pumped far from thermal equilibrium through photoexcitation with blue and ultraviolet light. Interpolating between the nonequilibrium experiments and fully ground-state equilibrium MD simulations, we discussed the inclusion of multi-reference CASSCF/CASPT2 electronic structure calculations into the excited-state charge distributions for the tryptophans in BSA, and the implications for understanding the microscopic origins of the observed phonon condensation phenomena. Inspired by a number of speakers in Workshop IV, including Claudia Filippi (“Targeting excited states with quantum Monte Carlo”), Lucas Wagner (“Obtaining compact representations of excited states from QMC, and using those excited states as data for low-energy models”), and Julia Westermayr (“Physically inspired machine learning for excited states”), we are pursuing possibilities for using higher-level electronic structure calculations—à la quantum embedding—and machine learning to advance our understanding of the atomistic dynamics of complex biomolecular systems.

Multiscale approaches: Complementary approaches pursued in the Excited and Open Quantum Systems and Field Theory working groups advanced our understanding of how electromagnetic excitations coordinate and organize complex systems. One intriguing observation that emerged from the discussion of excited states in classical MD simulations is the connection between divergent scales, from single tryptophan amino acids through the entire protein residue network and into the ionic and aqueous environment. These molecular reverberations from excited electronic states into the vibrational modes of the protein are manifested most prominently in fully nonequilibrium environments, and these low-frequency THz signatures arise from a complex interplay of dipole auto-correlations and cross-correlations among the various system components, including ions and water.

Work done at IPAM: Members of the Excited and Open Quantum Systems working group produced a series of lecture notes on restricted Hilbert spaces, ensemble DFT, and the nonequilibrium Lindblad and Mori-Zwanzig formalism. Two members of the group finished data analysis and manuscript preparation for a full-length research article, which they are currently preparing for submission to *PNAS USA.*

Outlook: The study of light-matter interactions has received considerable attention recently not only in the study of confined systems, e.g. cavity quantum electrodynamics, but also for large molecular systems relevant in quantum biology. The possibility of designing quantum cavities that leverage the interactions with specific modes of the quantized field to allow only specific collective excitations has opened new possibilities for materials design. Furthermore, the deployment of photoexcited systems from high-level quantum chemistry calculations into classical molecular dynamics for multiscale modeling across electronic and vibrational degrees of freedom presents new opportunities for integration into efficient, widely used simulation packages for a range of novel applications.

4. Embedding theories

Introduction: Embedding theories provide a natural framework to combine different computational scales within one large problem. Electronic-structure theory is rich in examples for which this embedding perspective is beneficial: lattice defects in crystals, molecules in close vicinity of a surface, or individual atoms near the active site in an enzyme—to name only a few examples. In any embedding theory, the full and commonly large problem is partitioned into one or multiple fragments of interest, and their respective environments. In this context it is natural to think of the individual fragments as (open) subsystems or impurities that interact with an environment that corresponds to the full system. Different flavors of embedding methods can treat the partitioning and/or interaction of the fragments with the environment differently. Yet the central theme of embedding methods is clear: provide a computational approach to different properties of the impurity that takes the interaction with the environment into account while maintaining a low computational cost. This idea is very general, and consequently encompasses a wide variety of approaches. In fact, any computational recipe that partitions a system and treats the individual parts at two different levels of theory is formally an embedding. For example, active-space calculations or frozen-core calculations can be seen as embedding approaches. A more concrete distinction of different embedding theories can be obtained by analyzing the individual models for the interaction between fragment and environment.

Density matrix embedding theory in the context of other embedding methods: While any combination of techniques constitutes an embedding, several rigorous formalisms have emerged that provide for exact feedback between the embedded system and its environment. Commonly, three approaches are distinguished in the literature: density functional embedding (e.g. density functional theory (DFT)-in-DFT and wavefunction (WF)-in-DFT), Green's function embedding (e.g. dynamical mean-field theory (DMFT), quantum defect embedding theory (QDET), and self-energy embedding theory (SEET)), and density matrix embedding (e.g. density matrix embedding theory (DMET) and variational embedding).

Throughout this long-term program, the focus was predominantly on different flavors of DMET, which was proposed in 2012 by Prof. Garnet K.-L. Chan. In DMET, the Schmidt decomposition is used to represent the entanglement between fragments and their environment through bath orbitals. The bath orbitals in DMET offer an advantage over density functional embedding, where interactions (e.g. van der Waals) between the fragment and environment may be neglected due to the approximate nature of the density functionals that are used in practice. Furthermore, unlike Green's function embedding methods, which require the calculation of frequency-dependent quantities, DMET can make use of the many existing *ab initio* quantum chemistry methods that have been previously developed for frequency-independent ground state calculations. Hence, DMET is a multi-level approach bridging different rungs in the electronic-structure ladder of accuracy. As such, multiple aspects of DMET, together with other embedding methods, were discussed during the first workshop at IPAM (*Workshop I: Multiscale Approaches in Quantum Mechanics*).

Interaction of fragment and environment: In DMET, the interaction between the individual fragments and their corresponding environment is taken into account through the so-called bath orbitals. The bath in DMET can be interpreted as the spillage of the fragment occupied orbitals into the environment. A more sophisticated motivation is through the Schmidt decomposition (SVD), which represents the entanglement between the fragment and the environment through a compressed environment—the bath. Mathematically, projecting the atomic orbitals of the fragment successively onto the occupied subspace and then onto the environment suffices. A more explicit matrix approach is by means of the cosine-sine decomposition of the 1-RDM. During the long-term program, alternative bath orbital constructions were discussed. Though generally important, the bath orbital construction becomes particularly important when extending DMET to allow for ensemble density matrices. In this general context, fractional occupancies of orbitals may result in the number of chosen bath orbitals to (potentially) be larger than the fragment size. Moreover, this adds an additional constraint since the number of bath orbitals needs to be judiciously chosen in order to avoid introducing a source of discontinuity on, for example, the potential energy surface. Since ensemble density matrices are in general not idempotent, the different bath orbital constructions are no-longer equivalent. Which particular bath orbital construction is beneficial in the context of ensemble densities is unclear at this time. It is moreover worthwhile to mention that including additional bath orbitals changes the key characteristic of DMET, namely, that the impurity systems are half-filled, i.e., the number of bath orbitals is (at most) equal to the number of fragment orbitals.

Chemical potential fitting: The individual impurities that are to be solved in DMET can be seens as “open” quantum systems. It is therefore highly likely that the fragment as well as bath orbitals are occupied. Since the global high-level 1-RDM consists merely of the fragment orbitals, the global number of electrons may not be preserved. The corresponding particle number constraint, i.e. the total number of electrons must be equal to the combined total number of electrons within each fragment, is enforced by employing the global chemical potential as a Lagrange multiplier. Throughout the high-level calculations, this chemical potential is fitted to enforce this constraint.

As opposed to conventional chemical potential fitting in electronic structure theory, it turns out that, in the context of DMET, this procedure is by no means straightforward. More precisely, the particle number in the individual fragments as a function of the chemical potential can in fact be discontinuous. This becomes particularly problematic when the targeted particle number lies within the discontinuity. Within this long-program, we have revealed that this issue already appears in the one- and two-dimensional Hubbard model with merely a few sites. Hence, it is worthwhile to recognize the chemical potential fitting problem within DMET as a research topic which requires careful considerations to determine the energy minimizing chemical potential.

Careful numerical considerations have revealed a sufficient condition for the discontinuity in the particle number, namely, the crossing of energy levels as a function of the chemical potential in the impurity problem. This finding allowed us to develop a chemical potential fitting procedure that formally remedies the discontinuity issue, christened the multi-state chemical potential fitting.

Work done at IPAM: A working group investigated the source of the discontinuities observed in the particle number in the impurity as a function of chemical potential. In the effort to address this, various methods were tested to detect and circumvent the numerical issues that arise from these discontinuities. The goal of this investigation is to form an autonomous, black-box, chemical potential fitting procedure in order to contribute to the development of a robust DMET code for molecules.

Outlook: Quantum embedding theories enable *ab initio* calculations for large systems within current computing capacities while preserving a high level of accuracy in chemically relevant regions. Open and important problems in DMET include the above outlined ambiguity of the bath orbital construction in the context of ensemble densities. Whether there exists a unique bath construction that is universally applicable to all flavors of DMET, or whether problem-specific bath constructions should be pursued remains unclear at this point. Moreover, a black-box practical implementation of the multi-state chemical potential fitting proposed above is yet to be implemented and incorporated into state-of-the-art DMET codes. This is strongly affected by the (quasi) degeneracy of the targeted states close to or at the discontinuity in particle number. Two possible remedies were proposed within the working group and are currently being numerically investigated. Furthermore, open problems in embedding theories more generally include a rigorous characterization of optimal fragmentation strategies and localization objectives.

5. Response functions in molecules and solids

Introduction: Describing the electronic structure, properties, and effective behavior of molecules, solids, and their interactions is at the heart of many applications such as reactions on surfaces, catalysis, self-assembly or molecular sensing. If considered separately however, two communities consider different methods: solids are regularly described with Density Functional Theory (DFT) using plane wave basis sets to describe the spatial electron density while molecules are often treated with Coupled Cluster (e.g. CCSD) or Quantum Monte Carlo (QMC) methods using atom-centered basis functions. This poses significant challenges for combined systems where computational feasibility and chemical accuracy needs to be balanced or for systems where the solution needs to be systematically improvable to settle discrepancies between reference methods or quantum chemistry and experiment. These problems persist even in the context of machine learning for quantum mechanics, since accurate training data is required.

Multi-level quantum methods: Different levels of theory that would be preferable for isolated components of a heterogeneous system (e.g. DFT and CCSD) typically cannot be derived from the same formal framework which requires a different formal analysis for each such method to obtain performance guarantees. Always choosing the highest level of theory for the whole system incurs a prohibitive computational cost, while choosing a lower level misses important aspects of the electronic structure. Consequently, there is a need for either machine learning models to capture the effective interactions or quantum chemistry methods that combine information from different levels of theory directly such as a QM/QM embedding or response functions of different order evaluated with different methods.

Systematic improvement: Some levels of theory belong to one family and form a series of strictly improving description of electronic structure (e.g. HF/MP2/CCSD). The basis sets for the description of the electron density or wavefunction, however, either improve very slowly albeit systematically (plane waves), are not well-established (wavelets) or approximate the complete basis set limit in an uncontrolled manner only (atomic bases). Given that plane waves are suitable for DFT but exceptionally expensive for higher level methods, a systematically improvable and computationally feasible expansion of the electron density is missing. This might be the reason why some high-level quantum chemistry methods (QMC and CCSD(T)) do not agree on non-covalent interaction energies for medium-sized molecules. This discrepancy reaches a chemically relevant magnitude, so systematically improvable descriptions remain desirable. In the context of model approximations, a systematic improvement is often feasible: in the case of machine learning models, more or better training data offers a clear pathway to improvement while quantum alchemy improves systematically upon inclusion of higher order derivative information.

Response functions such as the susceptibility or properties implicitly defined from response functions such as the polarizability require the derivative of the electronic structure result with respect to an external perturbation to be representable in the finite basis of the quantum chemistry calculation. This often is not the case, as basis sets are commonly optimized for the energy of a system rather than the electron density which converges more slowly with basis set quality. A good description of response functions is directly required for quantum alchemy as well as the analysis of chemical reactivity. Indirectly, machine learning models benefit from training data where responses or property gradients are consistent with numerical derivatives.

Work done at IPAM: An open challenge within the framework of non-covalent interactions is the numerical evaluation of non-local response functions of real atoms and molecules beyond model systems. During the programme, methods have been developed to evaluate these functions with arbitrary electronic structure methods which leads to benchmarks for different response models. This furthers the understanding of the real space behavior of response properties and facilitates the development of physically informed coarse-grained schemes needed for the modeling of larger systems. Currently only the non-correlated electric susceptibility is discussed in literature. The results from the programme now give access to different types of response functions and model them directly accounting for correlation effects. To improve the description of response functions embedding approaches have been investigated where the basis set of a variational electronic structure calculation gets extended locally to identify regions in space where the basis set is not flexible enough. The QM/QM embedding approach would help to keep computational requirements low, since only a small part of the electronic structure calculation needs to be updated to reflect this change. Potentially, this might open up a route to systematically improvable atomic basis sets which then effectively become a function of the system under consideration. This would also help to obtain more accurate response functions for quantum alchemy, since the current independent optimization of basis sets for each element introduces numerical artifacts. An ongoing effort initiated at IPAM is to calculate density response properties with accurate quantum Monte Carlo methods. Current efforts are with real-space variational Monte Carlo (VMC) and building toward implementations with diffusion Monte Carlo (DMC).

Benchmark Systems: Ideally, relevant example systems are accessible for many levels of theory. For response functions and polarizabilities, suitable systems are atoms, dimers (including systems with non-nuclear attractors) and small non-covalently interacting molecules. To improve the quality of atom-centered basis sets, an emphasis on small interacting molecules is clearly beneficial as it avoids overfitting towards atomic energies. This is expected to also improve the description of response functions and the generality of the resulting basis. Particularly suitable reference systems are those where high-level quantum chemistry methods disagree.

Outlook: Combining response functions such as alchemical derivatives may offer a pathway to combine and connect different levels of theory into one systematically improvable approximation of a finite volume of chemical space. These responses may find applications in materials design, reactivity analysis, and providing gradient information for machine learning amongst others. Considering alchemical derivatives in the development of new methods may help transitioning from the serial evaluation of systems towards methods that scale with the combinatorial complexity of chemical space. The new methods to obtain correlated non-local response functions may be beneficial to the theory of non-covalent interactions which, in large part, was developed using the concepts of these quantities without any means of calculating them directly. Based on these results, we expect to be able to design novel methods to treat non-covalent interactions on a clear physical basis.

6. Non-covalent Interactions

Introduction: Intermolecular interactions play a crucial role in determining the macroscopic properties of matter, particularly in large systems. The mathematical modeling of these interactions is, however, a problem of outstanding theoretical and computational difficulty: it involves, on the one hand, the resolution of deep mathematical and physical questions, and on the other hand, designing computationally tractable algorithms that can be applied to large-scale problems. Given the diverse backgrounds of the participants which included physicists, mathematicians, and computer scientists, this IPAM Long Program seemed particularly well-suited to address this problem. This led to several fruitful discussions and the beginning of interdisciplinary collaborations that will advance our qualitative understanding of the mechanisms of non-covalent interactions and extend the range of applicability of reliable numerical methods to large-scale systems. While the field of non-covalent interactions encompasses a broad range of phenomena, a major focus of the IPAM Long Program participants was on gaining a deeper understanding of so-called van der Waals (vdW) dispersion interactions.

The challenge of modeling dispersion in large-scale systems: Dispersion interactions arise from intrinsically quantum-mechanical fluctuations in the instantaneous electronic charge distribution and, thus, are inherently non-local, many-body, and long-range in nature. As a result, vdW dispersion is of particular importance in large (molecular) systems. In contrast to local and shorter-range (covalent) interactions, which are naturally confined to a limited number of nearest neighbors, dispersion forces involve extended, if not all, parts of the system. As such, they play a decisive role in shaping the global structure, dynamics, and response of large systems. Recent works have shown that failing to capture vdW dispersion in its full complexity has important implications. For example, adapting simplified formalisms can lead to qualitative failures in the description of non-covalently bound molecular complexes, molecular materials, or solvated biomolecules. The accurate description of dispersion interactions thus requires a sophisticated quantum-mechanical treatment. Only recently has the growth in computational resources made it possible to use such methods to showcase the role of non-locality and collectivity in a few, select examples. Routine application of such methods in large-scale simulations of systems such as amorphous molecular solids and solvated proteins, however, is still unfeasible. Part of this program therefore was an interdisciplinary effort to address the current lack of efficient, yet sufficiently accurate, methods for the description of vdW dispersion in large-scale systems.

Overview of the Methodology: The first step in this effort was to establish a common understanding of the vast range of existing methods within the language and context of the various backgrounds of the program participants. Due to the scope of the Long Program, special attention was paid to modeling dispersion interactions within the framework of density functional methods. Preliminary discussions focused on the fundamental shortcomings of the so-called “atoms-in-molecules” approaches. Most prominently these include an implicit assumption of locality, which naturally limits the capacity of these methods in describing intrinsically nonlocal phenomena as posed by vdW dispersion forces. To overcome this limitation, the participants focused on self-consistent interatomic models, which account for screening and polarization effects in a fully many-body fashion. One such method is the many-body dispersion (MBD) method. Over the course of the program, remaining limitations of this method, their solution — especially in the context of the fundamental formalism and its algorithmic implementation — as well as potential synergies with alternative approaches such as ML have been discussed.

A Second-Quantized Approach to Dispersion Interactions: Recent works on molecular complexes, vdW materials, and solvated proteins indicate that dispersion forces are highly delocalized and collective in extended systems. This in fact suggests adapting a formalism of collective, wavelike excitations rather than the current atom-centered formulation. Second quantization thereby provides a natural language for describing vdW dispersion in terms of delocalized collective excitations. As part of this IPAM Long Program, the MBD method has been formulated and discussed in the formalism of second quantization for analysis and understanding of vdW forces. In addition, it paves the way towards more reliable approximations: Currently used approximate models limit the description to the interaction of select numbers of atoms, which is in stark contrast to the inherent collectivity of vdW interactions. The framework of second-quantized MBD, on the other hand, allows one to focus on the most relevant contributions within a basis of collective excitations — potentially including all atoms — and thus capturing the delocalized, plasmonic nature of dispersion interactions.

Machine Learning Approaches: Despite tremendous progress and the success of machine learning (ML) potentials in the past years, data-driven predictions of long-range (dispersion) interactions still remain elusive. This is in part due to the complex quantum-mechanical nature of these interactions and in part due to the fact that dispersion interactions are crucially sensitive to the global configuration/conformation of the system, which presents an intractable sampling space for standard ML. To tackle this problem, a hybrid approach of ML and physical models has been proposed. Modern ML approaches to electronic structure theory are used to parametrize accurate physics-based methods to describe vdW dispersion without resorting to computationally demanding electronic structure calculations. During the IPAM Long Program, the first steps towards such a framework have been taken by combining accurate ML prediction of electron densities and the MBD method. Another promising approach envisaged and started during this Long Program is the development of ML models for an interaction matrix representation of the many-body vdW dispersion accessible through the second-quantized formulation of the MBD method. This allows us to learn effective all-to-all interactions within the quantum many-body state. Two particular advantages of this approach are the inherent account for many-body effects and the more detailed information content within the interaction matrix representation. This can considerably facilitate the learning procedure in comparison to the usual choice of the cumulative information contained in (total) energies or forces as used in existing methods.

Towards a Fundamental Understanding of Underlying Response Properties: Due to their practical importance, the development of theories of intermolecular interactions has largely been driven by applications leading to a zoo of approximate methods used by different scientific communities. Further advances in theoretical understanding of intermolecular interactions, however, require moving focus from specific applications to the general theory of electronic response functions of the interacting units, which are the natural physical quantities that underlie the description of intermolecular interactions. Revisiting the fundamental quantum physics of atomic and molecular response properties is therefore an important perspective.

The dipolar atomic polarizability is one of the most important among such response quantities since it determines the strength of induction and dispersion interactions. Previously unknown, fundamental connections between the characteristic size and the polarizability of the interacting units have recently been unveiled for real atoms as well as for quantum-mechanical model systems. Based on these findings, participants of the IPAM Long Program also focused on understanding the interplay between polarizability, the ground state geometry and the energetic spectrum of realistic systems. This focus led to a better understanding of non-local response functions, which was advanced by consistently formulating physical ideas in the language of distribution theory. In addition, thanks to constructive discussions with the participants working on the problem of basis sets for electronic structure calculations, the exact non-local distribution of correlated response functions for real atoms and molecules have, for the first time, been explicitly evaluated numerically. These findings are detailed in Section 5 (vide supra).

Outlook: In summary, this IPAM Long Program established a common understanding of the physical origins of non-covalent interactions and existing methods to describe them. To tackle the grand challenge of modeling vdW dispersion interactions in large-scale systems, the first steps towards novel, interdisciplinary approaches were taken. This includes hybrid approaches combining physical models and ML techniques as well as a second-quantization formalism for vdW dispersion and an hitherto underexplored approach based on understanding the connection between a system’s geometry and its non-local response function(s).

The developed methods thereby aim at both gaining a more fundamental understanding of these complex interactions in many-body systems as well as pushing the envelope of tractable length and time scales in simulations. Based on such firm theoretical understanding, systematic improvements of algorithmic implementations can ultimately allow us to routinely study realistic molecular systems and materials on a fully quantum-mechanical basis. In addition, this is expected to provide conceptual insights into the quantum mechanics of large-scale systems in general.

7. Machine learning for Quantum Mechanics

Introduction: Problems arising in quantum mechanics are by nature high-dimensional and nonlinear, requiring iterative (self-consistent) methods. In this regard, traditional numerical analysis techniques in conjunction with fully quantum mechanical simulations are computationally out of reach for chemical systems of significant size. Machine learning (ML) methods have proven to be a powerful technique for attacking these problems with high accuracy but at lower computational cost. Of particular interest are those that try to capture physically relevant information. The environment provided by the IPAM, bringing together researchers from diverse fields, including computer scientists, physicists, chemists and mathematicians, gave rise to fruitful discussions about the current range of applications, approaches, challenges and limitations of machine learning methods in quantum mechanics. The results of these discussions will be summarized below.

Local and global interactions with equal accuracy: Many properties of physical systems emerge from complex many-body interactions among atoms. Such properties are hard to capture within machine learning (ML) models, as they pose significant scalability challenges with respect to model complexity. For example, a set of pairwise interacting atoms already requires a quadratic number of interactions to be parametrized, which is why locality assumptions are typically introduced when modeling larger systems. As a consequence, non-local interactions are poorly or not at all described within local ML models, even if those interactions are contained within the reference *ab initio* data, leading to difficulty with generalization to unseen systems. One of the objectives of this program was to approach this challenge in a more controlled way and construct models that include all relevant interaction scales. This can be achieved in several ways, e.g. using deep learning-based dimensionality reduction techniques (such as transformer architectures) or through the development of more stable numerical solvers that allow efficient training of global models where all degrees of freedom remain fully correlated. Transformer models are neural network models designed around the attention mechanism, which attributes significance to each part of the input data based on the context, which is a natural fit for modeling interactions at arbitrary length scales among all components of an atomistic system. This topic was recurring in various workshops and working groups during this program, where commonly made efficiency trade-offs and approximations found in current models were extensively discussed.

Numerical challenges in highly correlated systems: Atomic correlations within large physical systems are notoriously hard to model, because they yield poorly conditioned interaction matrices that are the root cause for instabilities when solving inverse problems using numerical optimization algorithms. Inverse problems are ubiquitous: some ML algorithms that increasingly provide the basis for empirical physics models are trained by solving inverse problems (e.g. Gaussian processes), but they also appear in other areas, e.g. when solving the Schrödinger equation with certain model Hamiltonians (within the many-body dispersion (MBD) method, see Section 3). Combined closed-form and iterative solvers can ameliorate numerical stability issues by preconditioning the optimization problem using a low number of effective degrees of freedom, before iteratively converging the residual to zero. Finding effective preconditioners is an inherently interdisciplinary problem that crucially relies on application-specific expertise, which was uniquely brought together by this program.

Efficient SE(3) equivariant models: SE(3) equivariant machine learning models enable the preservation and propagation of orientation information in the form of representations in terms of spherical harmonics, providing more expressive representations of atomic environments and making it possible to predict complex tensorial properties that are equivariant with respect to the orientation of the atomistic system. While these models have been applied with great success in a number of areas within quantum physics, one major drawback to their invariant counterparts is that equivariant models require the manipulation of higher-order angular components in order to preserve the orientation, increasing the computational costs by a factor of $O(L^4)$ in the worst case scenario, where L is the highest angular degree used. The workshops and working groups during this long program enabled more involved discussions on ways to mitigate this additional cost, which can be done by implementing more efficient libraries for performing the tensor products necessary to manipulate the representations in terms of spherical harmonics, which require joint effort from the experts in different fields participating in the long program. Another approach that was explored throughout the long program was the restricting of the maximum angular degree, or the amount of coupling between the different representations, since high angular degrees and full coupling are not necessary for many practical scenarios.

Learning at different length scales: The majority of machine learning models for quantum mechanics assume locality (or short-rangedness) of all quantum mechanical interactions, meaning that beyond a certain cutoff radius the effects of these interactions are considered to be negligible. While this assumption enables the development of efficient models whose computational cost scales linearly with system size, it also prohibits the possibility of describing non-local interactions at longer length scales, which become more prominent with increasing system size.

Another popular approach when dealing with large system sizes is coarse graining, which combines neighboring groups of particles into fragments, which are then themselves treated as individual entities. Within the machine learning context, this would allow for the training of two or more models, depending on the levels of fragmentation hierarchy, with each model operating on a different length scale from individual atoms to the highest level of fragments.

Electronic structure learning: Most machine learning models for atomistic systems focus on training on a fixed set of properties, which requires training a new model whenever an application scenario requiring a new property arises. To mitigate this drawback, recently a series of approaches have been developed that attempt to learn the electronic structure (in the form of the wavefunction or electron density) as a central property of atomistic systems, by training the models on a set of electronic structure properties obtained from quantum mechanical calculations. This approach would enable the derivation of a wide range of other properties (such as energy, forces, polarizability, electric multipoles, excitation etc.) without the need for additional training, thus circumventing the cost of quantum mechanical calculation while retaining their generality and physical interpretability. This is complicated by the fact that electronic structure representations are intricate tensorial objects that transform non-trivially under molecular rotations, making them a challenging prediction target. However, electronic structure representations can be predicted accurately by using specialized equivariant architectures, which capture the intrinsic structure and symmetries of these objects. Given the ability to extract a wide range of properties from these models, at the IPAM we started exploring various possible applications that exploit machine-learned electronic structure, such as: speeding up many-body dispersion with polarizabilities from machine-learned electronic structure (see Section 3), efficiently generating infrared spectra using dipole moments derived from machine learned densities, and predicting spectroscopic properties by modeling excited-state electronic structures, among others.

Machine learning for quantum Monte Carlo simulations: Another promising application of machine learning techniques is the efficient representation of quantum states. The exponential scaling of the Hilbert space of quantum many-body systems makes highly accurate wavefunction-based *ab initio* calculations computationally prohibitive for large systems and entails the need for efficient approximations. This motivates the application of machine learning, which has proven to be a very effective tool for pattern recognition and feature extraction. In quantum Monte Carlo simulations neural network trial wavefunctions pushed the state-of-the-art by combining physical constraints with the expressiveness of deep neural networks. Being exact in principle (though statistical), the accuracy of quantum Monte Carlo methods is limited by the efficiency of modeling and optimizing the wavefunction ansatz only and machine learning-based quantum Monte Carlo simulations recently produced some of the most accurate results for small molecules and model Hamiltonians. Despite many similarities to other applications for the prediction of the electronic structure or learning quantum mechanical properties of molecules, machine learning for quantum Monte Carlo is different in that the quantum many-body problem is solved unsupervised, that is, without the requirement for training sets and the imposed limitations of the method used to obtain that data. Comparing neural network architectures, pointing out current limitations, and learning from the experiences of related research drives the progress in this emerging field and helps in establishing machine learning as an integral tool for quantum Monte Carlo calculations. Discussions at the IPAM addressed the efficient representations of many-body correlation, the implementation of physical priors, and the fast and reliable optimization of neural network wavefunctions. Furthermore, practitioners were stimulating the introduction of publicly available test sets for the benchmarking of new models and discussing the combination of codebases to facilitate direct comparison between the various existing architectures.

Implementing physical constraints into ML models and functionals: Transferability, that is, the ability of a model to make predictions on new data that it has not been trained on, is a major concern in applications of ML. Since chemical compound space is incredibly vast and diverse, developing ML models that are transferable is a central problem in chemistry and condensed matter physics. In this regard, a promising way to create highly transferable ML models which we have explored at the IPAM is to implement known physical constraints into ML models. One such thrust in this direction is imposing equivariance constraints (e.g. conservation of energy with translations/rotations of a molecule) in ML-derived potential energy landscapes. Another major thrust in this research direction was in the context of density functional theory (DFT). In DFT, ML can be used to create a data-driven functional, which maps the electron density of a chemical system to the energy of that system. In this regard, there are a number of known exact constraints that this functional must obey (e.g. the solution to the uniform electron gas). At the IPAM, we worked on ways to incorporate these known constraints into ML-generated DFT functionals to provide high accuracy and transferability.

Interpretability: Machine learning models, including those used in quantum mechanics, are inherently nonlinear. While this nonlinear nature endows the models with the expressiveness necessary to approximate complex functions, it also precludes a direct examination of how the inputs contributed to the output of the model. However, being able to examine and explain how a model came to its decision is essential for many real-world scenarios, helps guide model development, and in the context of quantum mechanics can be used to verify the 'physical plausibility' of the model and can even lead to novel insights about quantum mechanics. The field of explainable AI has produced a number of approaches for extracting such explanations from large classes of highly nonlinear machine learning methods, while an alternative view is that machine learning models should be self-explainable, i.e. the ability to provide interpretable results should be built into the model architecture itself. For example, attention based-graph neural networks naturally attribute a significance to the various interatomic interactions via the attention weights, while models that predict electronic structure properties such as wavefunctions and densities provide objects that can often be understood and interpreted by physicists and chemists. The regular interactions between physicists, chemists, mathematicians, and machine learning experts enabled by this long program spawned many discussions about the topic of interpretability and its various merits, resulting in an increased interest in applying and improving existing explainable models, as well as developing new ones.

Quantum Alchemy: Modeling properties in molecular space suffers from sparse training data. Including alchemical derivatives (i.e. derivatives w.r.t. changes in elemental composition) might help models to be more data efficient. Since alchemical derivatives are often cheaper to obtain than additional training points, this effect might be substantial. During the program, ideas were developed to potentially incorporate this information into machine models, e.g. by including derivatives with respect to nuclear charges in the loss function. While the cost of computing and optimizing over the gradients increases as higher-order derivatives are included, the hope is that we can significantly increase the overall generalizability of these models by including even just first- or second-order gradient information.

Mini-Workshop—Machine Learned Potentials: Machine learning potentials (MLPs) are still the most widespread and well-researched application of ML in quantum mechanics, since they enable very fast computation of potentials at a much higher level of accuracy (approaching DFT/CC) than classical empirical force fields. This in turn opens up the possibility to extract quantum mechanical observables from large scale molecular dynamics simulations which would be impossible to perform with standard quantum chemical methods such as DFT. During the first workshop hosted by the IPAM, "Multiscale Approaches in Quantum Mechanics", a number of researchers pointed out to some open challenges when using machine learned potentials/force fields for large-scale simulations, including but not limited to (1) capturing long distance interactions; (2) the connection between message-passing and many-body interactions; (3) the advantages and disadvantages of linear/multi-linear models compared to over-parameterized non-linear approaches such as neural networks; and (4) the physical motivations and relevance behind the way these models are constructed. This provided an opportunity to bring the community closer together, and we decided to organize a mini-workshop as a part of the long program with an explicit focus on machine learning potentials (primarily organized by Marcel Langer and Christoph Ortner). The workshop took place from the 9-12th of May, and consisted of a total of nine open-ended talks, with the duration of each talk ranging from 90-150 minutes. The talks covered various fundamental MLP topics in great detail, ranging from in-depth explanations of methods such as message-passing neural networks, atomic cluster expansion, equivariant models, attention mechanisms, to discussions about important challenges such as long-range interactions, error propagation and interpretability. The open-ended nature of the talks allowed for very long and interactive discussions by all participants, leading to an increased understanding of the various methods across the board, which served as an inspiration for future research topics and collaborations among a large number of the workshop participants.

Outlook: Machine learning will continue to be a vibrant area of research throughout the sciences. While ML provides a promising alternative to traditional approaches to solving the quantum many-body problem, a number of challenges still remain. Developing ML algorithms that are efficient, highly accurate, and transferable remains an active area of research in the field of quantum mechanics. In this regard, computational efficiency can be improved through coarse graining and local interaction models, and transferability can be achieved through enforcing known physical and symmetry constraints.