

The Potential of Neural Network Potentials

Published as part of ACS Physical Chemistry Au virtual special issue “Visions for the Future of Physical Chemistry in 2050”.

Timothy T. Duignan*



Cite This: ACS Phys. Chem Au 2024, 4, 232–241



Read Online

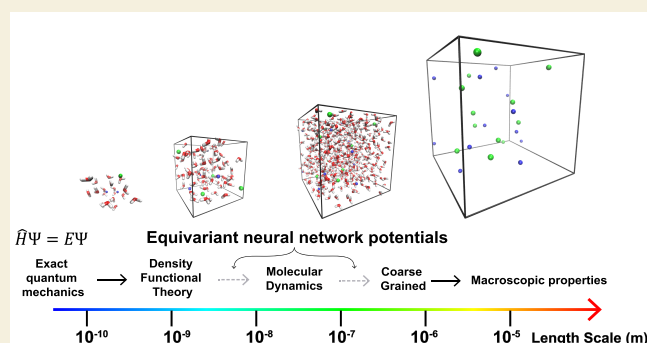
ACCESS |

Metrics & More

Article Recommendations

ABSTRACT: In the next half-century, physical chemistry will likely undergo a profound transformation, driven predominantly by the combination of recent advances in quantum chemistry and machine learning (ML). Specifically, equivariant neural network potentials (NNPs) are a breakthrough new tool that are already enabling us to simulate systems at the molecular scale with unprecedented accuracy and speed, relying on nothing but fundamental physical laws. The continued development of this approach will realize Paul Dirac’s 80-year-old vision of using quantum mechanics to unify physics with chemistry and providing invaluable tools for understanding materials science, biology, earth sciences, and beyond. The era of highly accurate and efficient first-principles molecular simulations will provide a wealth of training data that can be used to build automated computational methodologies, using tools such as diffusion models, for the design and optimization of systems at the molecular scale. Large language models (LLMs) will also evolve into increasingly indispensable tools for literature review, coding, idea generation, and scientific writing.

KEYWORDS: Molecular simulation, molecular dynamics, machine learning, artificial intelligence, coarse-graining, large language models, diffusion models



INTRODUCTION

Over the last year, a surge of excitement has arisen around the potentially transformative role of machine learning (ML) or artificial intelligence in many areas of technology, science, and society. This has been driven by the remarkable performance of large language models (LLMs)¹ and image generation diffusion models.²

Given the excitement around these technologies, there is a natural tendency to either assume they will profoundly transform every field or dismiss them as “overhyped”.

A better approach is to be excited about the potential of these technologies while also thinking through as carefully as possible the potential obstructions that may arise to limit their potential, with the goal of identifying the most promising approaches and applications where they are likely to be most useful.

While the recent excitement around LLMs and diffusion models is justified, the potential for ML in science has long been recognized. A pioneering example is the demonstration in 2007 by Behler and Parrinello³ that neural network potentials (NNPs), originally developed in the 1990s,^{4,5} could be used to massively accelerate the simulation of liquids by learning the underlying potential energy surface from quantum chemistry calculations.

More recently, AlphaFold2 (AF2)^{6,7} has dramatically demonstrated the ability of ML to predict the structure of folded proteins with remarkable accuracy.

In this article, I argue that neural network potentials, AF2, diffusion models, and even LLMs to some extent can all be understood as examples of one fundamental idea: learning and sampling from an underlying energy surface. The recursive application of this fundamental idea will have profound implications for physical chemistry and beyond over the next half-century.

DIRAC’S DREAM


Using quantum mechanics to predictively understand biology and chemistry has been a dream of physicists for nearly a century since Paul Dirac’s famous statement:

Received: January 18, 2024

Revised: March 4, 2024

Accepted: March 5, 2024

Published: March 21, 2024



*The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations that are much too complicated to be soluble. It therefore becomes desirable that approximate practical methods of applying quantum mechanics should be developed, which can lead to an explanation of the main features of complex atomic systems without too much computation.*⁸

Physical chemistry is the natural discipline to home this dream, and much progress has been made on making it a reality. For example, accurate quantum chemistry algorithms, such as density functional theory (DFT)⁹ and coupled-cluster (CC) with single, double, and perturbative triple excitations (CCSD-(T)),¹⁰ and advanced simulation techniques have made it possible to build highly accurate models of water¹¹ and resolve challenging scientific questions¹² starting from nothing but Schrödinger's equation. This pioneering work convincingly demonstrates the feasibility of achieving Dirac's dream. Now the key scientific challenge becomes making this procedure efficient, automated, generalizable, and scalable. This is a task perfectly suited to ML. Achieving this dream, depicted in Figure 1, would



Figure 1. DALL-E's depiction of the potential of neural network potentials.

be truly transformative, enabling the predictive understanding of the properties of a vast range of critically important systems throughout materials science, chemistry, biology, and beyond.

■ NEURAL NETWORK POTENTIALS

The most promising tool for doing this is neural network potentials (NNPs), which are a class of machine learning interatomic potentials (MLIPs). These tools are rapidly gaining importance and recognition due to their flexibility and efficacy.^{13–18}

The concept of NNPs is straightforward yet innovative. Instead of relying on traditional methods, such as specifying Lennard-Jones potentials and bonded terms, to define the potential energy surface (PES) as a function of atomic type and coordinates, NNPs employ more adaptable functions with a

larger set of parameters. These parameters are optimized using machine learning algorithms facilitated by automatic differentiation. To avoid the need for large experimental data sets, these models are trained on data from quantum chemistry, predominantly at the DFT level, though increasingly higher levels of theory are being utilized.¹⁹ Traditional molecular dynamics (MD) algorithms can then be used to simulate the direct physical behaviors of gases, liquids, and solids using the trained NNPs.

While NNPs have been under development for some time,¹⁴ recent advancements in ML have significantly enhanced their practical utility. One such critical advancement is the incorporation of equivariance, which encodes the spatial symmetries of 3D space into the neural network.^{20–23} This feature, also pivotal in the success of AlphaFold2,²⁴ ensures that the neural network's features rotate in accordance with the rotations of the input coordinates, significantly enhancing model accuracy and reliability. A recent review provides a comprehensive and clear explanation.²⁵

The success of equivariant models is very intuitive, as they use a natural extension of the traditional concepts of intermolecular modeling. More specifically, they use tensors to describe the atomic properties and tensor products for modeling the interactions. This approach mirrors the long-standing practice in physics, where molecular interactions have been described using the multipole expansion, i.e., using monopoles, dipoles, and quadrupoles, which are tensors, interacting via tensor products.

Additional exciting tools such as active and delta learning as well as improved quantum chemistry algorithms have also significantly improved the generalizability and usefulness of NNPs.²⁶ An additional avenue is the development of foundation models, where increasingly larger and more diverse/general training data sets are used to improve a model's generalizability.^{27–30}

These more recent improvements have stepped up NNPs from being an occasionally useful tool to a potentially transformative one. As some recent perspective articles state, this field is undergoing “remarkable/breathtaking progress,” achieving results that would be “unimaginable with conventional methods.”^{31,32}

The potential of neural network potentials is that they are a tool that provides access to a new spatial and temporal scale that has not yet been observed accurately. Such tool-driven revolutions are a recurring process in the history of science and are inevitably followed by a string of rapid and surprising discoveries. The telescope, the microscope, and the particle collider are famous examples. I believe that NNPs have the potential to approach a comparable level of significance. Initial examples of this are already beginning to emerge. For instance, important new insights into catalysis at high temperatures have already been discovered using NNP-MD simulations.³³ Additionally, NNP-MD is apparently now the state of the art for protein–ligand binding affinity predictions.³⁴

■ BREAKING THE ACCURACY–EFFICIENCY TRADE-OFF

The promise of NNPs lies in their potential to transcend the long-standing trade-off between accuracy and efficiency in simulations. Specifically, they offer the possibility of achieving quantum mechanics (QM)-level accuracy at the computational cost of classical MD. This breakthrough implies that this technique will become increasingly prevalent.

An important recent advance on this front has been the development of density corrected DFT methods,^{9,35} where the computation of the energy and electron density are separated to achieve a better balance of accuracy. These methods have been demonstrated to provide an improved description of aqueous solutions^{35–39} while maintaining affordable computational requirements. They are therefore well-suited to providing training data for NNPs. We have recently shown that accurate and efficient prediction of key electrolyte solution properties is possible by combining this method with an equivariant NNP.^{40,41}

REMAINING CHALLENGES

Although rapid progress is being made, several challenges remain toward fully realizing the potential of NNPs. One example accounts for long-range interactions. These are known to play an important role in many liquid phase phenomena.^{42,43} Multiple promising solutions are being explored.^{44–52} Another is the generation of sufficiently diverse training data sets, where active learning plays a crucial role. In active learning, the uncertainty within the model's predictions is monitored, and data points with higher uncertainty are added to the training set to enhance its accuracy.^{53–55} The use of structural databases generated at lower levels of theory is another promising option.^{17,56}

The accuracy of the underlying quantum chemical method used to generate the training data is also an issue, but there is significant exciting progress on this front too, as discussed below.

Speed and scalability also present significant hurdles. Equivariant NNPs, while powerful, tend to be slower than classical MD and standard MLIPs. Although they can handle large systems,^{57,58} this requires substantial computational resources. One solution is using equivariant NNPs to generate data for training faster non-equivariant architectures. There are many promising architectures from the more general family of MLIPs such as permutation invariant polynomials (PIPs) and Gaussian approximation potentials (GAPs), which generally appear to have similar performance to standard Behler–Parrinello-type NNPs and may be able to provide speed advantages.^{59,60} Data augmentation may be able to avoid the need for equivariant models as well.⁶¹ Another possibility is optimizing classical MD potentials against quantum data using automatic differentiation⁶² or algorithmic improvements.^{63,64}

RECURSIVE COARSE-GRAINING

Despite the promise of current approaches to improving speed, there are inherent limitations to any simulation that attempts to keep track of the position and motion of every single atom in a system. A more transformative solution lies in coarse-graining, which involves focusing on only the essential degrees of freedom of a system while neglecting the others. This is the essence of physical modeling at all scales. For example, in modeling a macroscopic object, only the center of mass of each atom is considered. The core motivation is that most of the system's information is often irrelevant for predicting its properties.

In physical chemistry, details such as the exact positions of every solvent molecule are usually not critical for understanding the properties of interest. Thus, continuum or implicit solvent models are used, which are a form of coarse-graining. Currently, these models are developed in a somewhat empirical manner, and a systematic and rigorous approach is needed. ML, with its fundamental focus on finding lower-dimensional representa-

tions of high-dimensional data sets, offers an excellent set of tools to address this challenge.^{65–68}

Remarkably, it is precisely the same technique that provides one possible solution: equivariant neural network potentials, as has recently been demonstrated.^{41,69–71} It is not surprising in retrospect that these tools are suitable for this task. The goal of coarse-graining is to find the potential of mean force (PMF) for a reduced set of degrees of freedom. This is just an energy surface, technically a free energy, and so this is a task well-suited for NNPs. Here NNPs have the key advantage over fitting a classical force field, as the free energy surfaces modeled by NNPs can capture a more complex-shaped landscape than the typical bonded or Lennard-Jones interactions used in a classical force field. In fact, fitting an NNP to the full PES is a form of coarse-graining itself. It is just the electronic positions that are the degrees of freedom that are being coarse-grained out. Protein folding models such as AF2 and RFDiffusion can actually be interpreted as an NNP that learns a coarse-grained free energy, as discussed below.

So the promising conclusion of this is that all atom equivariant NNPs trained on QM data can be used to recursively train coarse-grained NNPs.⁴¹ This could enable large-scale biological simulations starting solely from first principles. This process, illustrated in Figure 2, may have a transformative impact on




Figure 2. Depiction of a method of connecting the microscopic scale to the macroscopic scale via the recursive application of equivariant neural network potentials.

biology and materials science if the remaining challenges can be overcome and may eventually supplant tools like AlphaFold2. We have demonstrated the feasibility of this process of recursively training a coarse-grained equivariant NNP on its own output in a recent preprint for simple electrolyte solutions.⁴¹

Currently, coarse-graining with an NNP is applied to relatively small-scale classical systems, such as simplifying water to a single site⁷² or for proteins.^{67,69,70,73} However, they could potentially be iteratively applied to increasingly larger scales, perhaps one day encompassing complexes of multiple proteins, cell membranes, etc.

While the core principle of automatically creating reduced-dimensional representations will likely remain constant, advancements may be required in determining which features to retain at a higher resolution, like binding sites, and which to coarse-grain; here again, ML will be a critical tool.⁷⁴

This methodology may even be a general solution to the long-standing and central challenge of molecular simulation, which is how to connect scales, i.e., how to use information from the molecular scale to inform and understand meso- and macroscopic-scale phenomena. It is also possible to reproduce fine-grained representation from the coarse-grained one.⁷⁵

One importance caveat is that coarse-graining does distort the underlying dynamics, but solutions to this problem are being developed.^{76,77}

■ SIMULATING LARGE-SCALE FROM THE MOLECULAR SCALE

The goal of this methodology is both ambitious and thrilling. It opens the possibility of simulating large-scale complex systems starting solely from Schrödinger's equation. We can envisage a future where, through a simple web interface, we input a system of atoms and obtain a highly accurate depiction of their behavior at any desired scale. This would be a more general version of the way AF2 now enables the rapid prediction of the folded structure of a protein with a simple web interface.

Such a tool would be transformative for many areas of science. For a great many scientific problems, the scales at which the most critical and interesting processes occur are too small to observe experimentally either with direct microscopy or with more indirect methods, such as spectroscopy or X-ray diffraction. As a result, we are often left to rely on guess work, intuition, and indirect evidence to infer what is actually happening on the molecular scale. In other words, when it comes to the molecular scale, we are essentially operating in the dark. Accurate and efficient molecular-scale simulations could act as a computational microscope, shedding light on this critically important scale and opening a wealth of new otherwise hidden information by enabling us to simply directly observe the key phenomena that occur at this scale. The application areas span most of biology and much of chemical engineering, such as catalysis and electrochemical energy storage.

Additionally, we can use these simulations in combination with statistical mechanics to calculate key experimental properties. For example, free energies or activities can be directly estimated from molecular-scale simulations using tools such as the Kirkwood–Buff theory. Additionally, diffusivities and other kinetic properties can also be determined. These values can then be fed as parameters into simulations of larger-scale systems, where these parameters are a key source of uncertainty. For example, the activity of ions in solution is a key challenge for predicting solubilities and chemical equilibria. Similarly, climate simulations are currently limited by the lack of information about molecular-scale processes, such as the stability and formation of aerosols. Molecular simulation models that can quickly predict the distribution of aerosol particles could be connected to larger-scale climate simulation models to improve their reliability. These tools therefore have the potential to have an impact on many different areas of science across many different scales of phenomena wherever physical-based modeling is important. Additionally, in combination with automated coarse-grained systems, these can potentially be scaled to simulate systems of ever-increasing size by selectively ignoring aspects of the system that are not considered practically important.

■ QUANTUM CHEMISTRY ACCURACY

A critical limitation of NNPs is the current accuracy of methods for solving Schrödinger's equation to generate data for training NNPs. However, the field of quantum chemistry is witnessing substantial progress with exciting developments like density corrected DFT,³⁵ ensemble DFT,⁷⁸ DLPNO–CCSD(T),⁷⁹ large-scale random phase approximation (RPA),⁸⁰ and the density matrix renormalization group.⁸¹ These innovations promise more precise descriptions of molecular-scale systems. Additionally, ML algorithms could also have potentially important applications here by using neural network ansatzes

for wave function representation, for example,⁸² or for reducing the cost of traditional methods.^{83,84}

While DFT functional accuracy is still an issue, a key advantage is that for a great many systems of practical interest, methods such as CCSD(T) can generate very high-quality data on small systems. This means that DFT functionals can be carefully validated and fine-tuned for specific tasks. This means that the problem of DFT functional accuracy is for many cases mainly a practical problem rather than a fundamental scientific challenge. Exciting new tools for automated development of new DFT functionals trained on data from higher-quality quantum chemistry methods is another promising recent development.⁸⁵

However, for some materials, we still lack good benchmark methods. For example, highly multireference states, where fundamental new tools are required. Multireference methods such as complete active space self-consistent field (CASSCF) can be used, but they come with their own set of challenges and limitations, particularly in terms of computational cost and complexity.

■ COMPUTATIONAL HARDWARE

Another potential avenue for advancement lies in the development of new, powerful hardware specifically designed for molecular simulation and quantum chemistry calculations. While quantum computers are often cited in this context, it is debatable whether using qubits will provide a practically relevant quantum chemistry acceleration.^{86,87} Other burgeoning technologies, such as thermodynamic computing, may prove useful instead. Additionally, the ongoing rapid development of new computational hardware for general computing and ML, such as large-memory graphics processing units (GPUs), should be sufficient to provide significant improvements in quantum chemistry and molecular simulation using existing algorithms.

The development of new computing architectures relies on components that are becoming increasingly small and will continue to approach the molecular scale. The design of these components could potentially be significantly advanced through highly accurate molecular simulations, which could be used to predict and optimize important properties such as heat generation, etc. Improvements in molecular simulation enabled by NNPs could therefore lead to a feedback process in which they enable improvements in computing.

■ DIFFUSION MODELS

Accurate molecular simulation is a powerful tool for understanding and predicting the properties of various systems. However, when it comes to generating and designing new systems, an additional new tool from ML is increasingly coming into play: diffusion models. These models, though initially complex to grasp, especially as often presented, are fundamentally about learning a generalized version of a potential energy surface.

To conceptualize diffusion models from a physical point of view, consider the pixel values to represent atomic coordinates. Then, just as certain coordinates have a high probability of being observed in an equilibrium simulation, certain pixel values have a high probability of having certain properties, i.e., being a face.

To generate these high-probability coordinates/pixel values, we evolve the coordinates/pixel values according to equations of motion with a force field, which is learned from the data via a denoising process. In fact, a diffusion model can be trained on molecular coordinates from an equilibrium distribution, which,

in turn, reproduces the true underlying force field responsible for that distribution.⁸⁸ Diffusion models then can be considered again as a special case of an NNP, where they are learning a coarse-grained free energy surface. An intriguing aspect of these models is that they can include “forces” for properties other than positions, such as atomic numbers, allowing them to sample from a variety of atomic types. This flexibility enables targeted sampling from specific regions of the probability distribution, akin to biasing MD simulations toward areas of the phase space of particular interest.^{89,90} Technically a time-dependent force is used to improve the convergence of the sampling, but this is a technicality.

This methodology is increasingly promising for the generation and design of new materials and molecules. It is already proving capable of protein design, leveraging databases like the Protein Data Bank (PDB) to identify new specific folded protein structures.⁸⁹

■ ALPHAFOLD2

One potential challenge to the importance of NNPs is AF2. AF2 represents a significant breakthrough in protein structure prediction. Its capability allows anyone to predict the atomically resolved structure of a protein, with accuracy now comparable to that obtained from nuclear magnetic resonance (NMR) techniques,⁹¹ marking a paradigm shift in structural biology. While this is not traditionally considered a core problem of physical chemistry, there are many closely related problems, i.e., predicting the atomic structure of a chemical system in thermodynamic equilibrium. Most significantly, AF2 does not explicitly rely on physical notions of energy and MD. An intriguing question then arises as to whether this approach could be used to revolutionize biology and chemistry without the need for physics-based modeling.

This is highly unlikely, in my opinion. The key to AF2's success is the extensive, high-quality Protein Data Bank (PDB).⁹² This immediately highlights a challenge to generalizing this approach. The PDB contains high-quality atomic structures of over 100 000 proteins with each one taking significant investment to obtain. This took many decades to build. There are vanishingly few problems of practical scientific interest, especially in physical chemistry, where these kinds of extremely large, high-quality databases exist. This becomes evident as soon as we leave the region of AF2's success.

For instance, AF2 cannot predict the folding pathway or kinetics of a protein due to its reliance on the static folded structures found in the database.⁹³ Furthermore, the database lacks information on the distributions of many ions near the protein surfaces, limiting AF2's ability to predict properties influenced by electrolyte concentrations, such as protein agglomeration, ion conductivities, etc.

Additionally, AF2 has a deep conceptual relationship to NNPs. It has been demonstrated that AF2 can be interpreted as using a learned energy landscape to optimize the protein structure;⁹⁴ the main difference is that AF2 is trained on experimental obtained structures rather than quantum chemical data, but experimental structures still contain important information about the energy landscape: they must lie on its minima. As AF2 uses gradient descent to optimize the protein structures, it must have learned an energy landscape with the same energy minima as the true protein energy landscape. The true protein energy landscape can be rigorously described as the coarse-grained free energy surface where the solvent degrees of freedom have been integrated (or marginalized) out. This

energy surface may look quite different from the true free energy surface far from a minimum, but this is not an issue for obtaining minimum energy structures, although it is for dynamic information.

For diffusion models trained on the PDB, the connection is even clearer. The score learned by the model is equivalent to the gradient of a coarse-grained free energy. If one could train the diffusion model on the true equilibrium distribution of states occupied by the protein, then one would learn the true coarse-grained free energy. The use of PDB training data means that one is effectively training on coordinates extracted at 0 K; so the off-equilibrium distribution will be wrong, but it will still have the correct minima. In fact, Langevin dynamics are used in diffusion models, which is a common algorithm to perform molecular simulations. Thus, something like RFdiffusion⁸⁹ is actually a coarse-grained MD simulation just on an approximation of the true coarse-grained free energy surface.

■ THE ONLY ALTERNATIVE?

The dependence of AF2 on extensive high-quality experimental data sets indicates that the same strategy will not work for predicting a broader range of important phenomena. This situation has a parallel to the development of AlphaGo, an ML algorithm that can play the board game Go, which was trained on data from human games accumulated over many decades. To improve the performance beyond what was learnable from human games, AlphaZero was developed, which does not require any human-generated training data and learns through self-play based only on the rules of the game.⁹⁵ Similar recent breakthroughs in mathematical problem solving use a similar strategy.⁹⁶

A natural question then arises: can we use this same strategy for solving scientific problems to avoid the need for large, hand-curated experimental databases? Doing so would require one to know the fundamental rules of the game that nature obeys. Of course, this is exactly what quantum mechanics is. It can in principle predict the behavior of all biological and chemical systems as captured by Dirac's dream. NNPs are therefore the best option for using ML to have a broad and significant impact on physical chemistry and beyond.

■ ROBOTICS AND SELF-DRIVING LABORATORIES

Laboratories will also become increasingly automated as robotics continues to improve, reducing the need for humans in the loop experimentation and enabling automated active learning feedback where errors in computational methods are identified and rectified through iterative exploration of regions of high uncertainty.

■ END-TO-END DIFFERENTIABILITY

Another critical advancement, which was also key to the success of AF2, is the concept of end-to-end differentiability. Essentially, this means that it should be possible to take derivatives of any parameter used in modeling a system with respect to the final loss function. This approach is exemplified by graph neural networks, which many modern NNP architectures successfully use. The idea is that rather than imposing arbitrary descriptors of the atomic structure, a flexible and general function can be used with many adjustable parameters that are then optimized using stochastic gradient descent just as the other parameters in a neural network are fitted.⁹⁷

End-to-end differentiability is also showing promise in applications to simulations themselves where gradients can be taken through simulations, allowing the automatic optimization of parameters with respect to output properties of molecular simulations. For example, this can be used to fit interatomic potentials to experimentally obtainable structural data.⁹⁸ It is also an important component of modern generative models for drug discovery.⁹⁹ Finally, this idea is also being applied to DFT functional development, where all of the key parameters used to define a DFT functional are treated as differentiable parameters that can be optimized with respect to some loss.⁸⁵

■ LARGE LANGUAGE MODELS

An obvious impact of ML in physical chemistry is through large language models (LLMs), e.g., ChatGPT, which have dramatically expanded their capabilities recently. These models can assist in various academic activities, including writing, brainstorming, literature searches, education, and even language translation for non-native English speakers. Their utility is expected to increase as they become more accurate and less prone to producing erroneous outputs (“hallucinations”). Recent articles provide a detailed demonstration of the potential of LLMs in chemistry more generally.^{1,100}

LLMs will become increasingly integral to scientific meetings and conferences, offering real-time notes, suggestions, and relevant literature based on the extensive training data encoded in their weights. They will play a crucial role in reviewing and improving the scientific literature, making research more accessible and efficient. Moreover, their capability in condensing vast amounts of data will likely foster novel connections and insights in research.

Furthermore, LLMs will be instrumental in programming, aiding researchers with limited coding proficiency in translating their ideas into executable code or converting code across different programming languages. This will significantly enhance scientific productivity and democratize computational tools, making cutting-edge research accessible to a broader spectrum of scientists.

LLMs will help automate the generation of MD trajectories and the building and curation of large databases. LLMs will not only suggest improvements for optimizing system performance but may also autonomously manage simulations, drawing from vast data sets to propose novel research directions.

There is also an intriguing connection between NNPs and LLMs. LLMs use a Boltzmann distribution, via the softmax activation function, to compute probabilities from the output of the transformer during next word prediction. They can therefore be interpreted as predicting an energy for each token in the vocabulary as a function of the previous words in the sentence, i.e., as a series of NNPs. While this is a more tenuous connection than for diffusion models, it hints at the deep connections between statistical mechanics and deep learning, which remain to be fully uncovered.

■ LONGER TERM

As we look beyond the next decade, two potential trajectories emerge. One radical possibility is that LLMs continue to improve at the current exponential pace, potentially surpassing human capabilities in 50 years, leading to a future dominated by ML-driven science. Alternatively, these tools might reach a plateau, maintaining their status as invaluable aids but still

leaving room for significant human contribution. This latter scenario presents an intriguing landscape for future speculation.

The likelihood of LLMs reaching a developmental plateau arises from three considerations. First, they are trained on Internet text, which, while extensive, inherently limits their capacity to generate novel knowledge. Second, the issue of “hallucinations” or generating false information remains a challenge. While scaling up models and refined training methods (such as reinforcement learning from human feedback) may mitigate this, this may be a more fundamental and challenging problem to resolve. The third reason that LLMs might plateau in performance may be the excessive energy demands of training ever larger models. Here again, advances in materials science to design new, more efficient forms of computing will be key, and improved molecular simulation enabled by NNPs will be a critically useful tool for this.

■ A MORE RADICAL FUTURE

One potential avenue to address the issues of hallucination and generalizing beyond the training data is to enable LLMs to verify or validate their generated knowledge. For example, in the area of physical chemistry, this could be enabled by augmenting an LLM by providing it access to molecular simulation tools such as NNP-MD and quantum chemistry. It could then automatically go about generating the necessary data and validating hypotheses it has generated autonomously. The likelihood of such a transformative leap in ML capabilities is extremely hard to ascertain, but if it does happen and continues to grow exponentially, it will be such a significant shift across all spheres of human society that speculation is only of limited use. In my personal view, while a temporary plateau in ML development might persist for decades, it seems improbable that it will extend for the next 50 years, and so I would expect at some point to witness profound changes in the way science and society operate.

This evolution raises profound and challenging questions both scientifically and ethically. A particularly sobering prospect is the potential redundancy of human involvement in scientific discovery courtesy of highly efficient LLMs or their successors. In such a scenario, the human role might be relegated to determining the tasks we would like ML to solve rather than engaging in direct scientific inquiry.

In the field of physical chemistry, many fundamental questions bear immense practical significance. It might be necessary to leverage ML’s capabilities to address these. However, we could also consider one day preserving certain areas within science as “ML-free zones”, allowing human scientists to explore them independently. Alternatively, we might use ML to develop technologies to enhance our quality of life while deliberately avoiding full comprehension of their workings to preserve a sense of purpose and discovery for ourselves.

While the prospect of ML solving complex scientific problems might initially seem disheartening, especially considering the potential obsolescence of human-led research, it is important to recognize the brighter side. In a world where ML resolves most material challenges, we could exist in a state of near-universal abundance. This shift might mitigate any disappointment stemming from the reduced role of human discovery. Moreover, even in such a world, humans could still find purpose in understanding and replicating ML-derived solutions, thus continuing the quest for knowledge and innovation.

CONCLUSIONS

ML is poised to revolutionize physical chemistry, reshaping our approach to solving many long-standing problems. A pivotal advancement facilitated by ML combined with state-of-the-art quantum chemistry is the ability to conduct molecular-scale simulations with unprecedented accuracy and speed. While there are still challenges to be overcome, this is already beginning with equivariant neural network potentials (NNPs) trained on extensive data sets derived from solutions of Schrödinger's equation. To access increasingly larger scales, these models will be recursively trained to automatically generate good coarse-grained descriptions, i.e., reduced dimension representations, thereby solving the long-standing challenge of connecting the microscopic and macroscopic scales and dramatically widening the scope of applications where first-principles physical simulation plays a critical role.

LLMs will improve efficiency and knowledge discovery. End-to-end differentiable simulations will efficiently compute the thermodynamic and kinetic properties of diverse systems. Diffusion models will be used to automate the generation of new potential candidate molecules and materials with particular properties.

Eventually, upon finalizing a design, it will be synthesized in automated laboratories, robotically produced, and dispatched for testing and further optimization based on experimental data. This vision heralds a new era in our understanding and manipulation of chemical systems, offering boundless opportunities for discovery and innovation.

Finally, it is worth considering the possibility of continued exponential advancements in ML. The implications of such a development are challenging to predict, but they could potentially reduce the field of physical chemistry to a discipline where our primary role is to ask LLMs to teach us complex concepts. The timeline for such a shift, whether it spans decades or centuries, remains an intriguing and open question.

AUTHOR INFORMATION

Corresponding Author

Timothy T. Duignan – Queensland Micro- and Nanotechnology Centre, Australia, Griffith University, Nathan, Queensland 4111, Australia; School of Chemical Engineering, University of Queensland, Brisbane, Queensland 4072, Australia; orcid.org/0000-0003-3772-8057; Email: tim@duignan.net

Complete contact information is available at:
<https://pubs.acs.org/10.1021/acsphyschemau.4c00004>

Notes

The author declares no competing financial interest.

ACKNOWLEDGMENTS

T.T.D. was supported by an Australian Research Council (ARC) Discovery Project (DP200102573) and a DECRA Fellowship (DE200100794).

REFERENCES

- (1) Boiko, D. A.; MacKnight, R.; Kline, B.; Gomes, G. Autonomous Chemical Research with Large Language Models. *Nature* **2023**, *624* (7992), 570–578.
- (2) Ho, J.; Jain, A.; Abbeel, P. Denoising Diffusion Probabilistic Models. *arXiv* **2020**, 1 DOI: [10.48550/arXiv.2006.11239](https://doi.org/10.48550/arXiv.2006.11239).
- (3) Behler, J.; Parrinello, M. Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces. *Phys. Rev. Lett.* **2007**, *98* (14), No. 146401.
- (4) Sumpter, B. G.; Noid, D. W. Potential Energy Surfaces for Macromolecules. A Neural Network Technique. *Chem. Phys. Lett.* **1992**, *192* (5–6), 455–462.
- (5) Blank, T. B.; Brown, S. D.; Calhoun, A. W.; Doren, D. J. Neural Network Models of Potential Energy Surfaces. *J. Chem. Phys.* **1995**, *103* (10), 4129–4137.
- (6) Jumper, J.; Evans, R.; Pritzel, A.; Green, T.; Figurnov, M.; Ronneberger, O.; Tunyasuvunakool, K.; Bates, R.; Židek, A.; Potapenko, A.; Bridgland, A.; Meyer, C.; Kohl, S. A. A.; Ballard, A. J.; Cowie, A.; Romera-Paredes, B.; Nikolov, S.; Jain, R.; Adler, J.; Back, T.; Petersen, S.; Reiman, D.; Clancy, E.; Zielinski, M.; Steinegger, M.; Pacholska, M.; Berghammer, T.; Bodenstein, S.; Silver, D.; Vinyals, O.; Senior, A. W.; Kavukcuoglu, K.; Kohli, P.; Hassabis, D. Highly Accurate Protein Structure Prediction with AlphaFold. *Nature* **2021**, *596* (7873), 583–589.
- (7) Jones, D. T.; Thornton, J. M. The Impact of AlphaFold2 One Year On. *Nat. Methods* **2022**, *19*, 11–26.
- (8) Dirac, P. A. M. Quantum Mechanics of Many-Electron Systems. In *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*; Royal Society, 1929; Vol. 123, No. 792, p 714–733. .
- (9) Sim, E.; Song, S.; Vuckovic, S.; Burke, K. Improving Results by Improving Densities: Density-Corrected Density Functional Theory. *J. Am. Chem. Soc.* **2022**, *144* (15), 6625–6639.
- (10) Raghavachari, K.; Trucks, G. W.; Pople, J. A.; Head-Gordon, M. A Fifth-Order Perturbation Comparison of Electron Correlation Theories. *Chem. Phys. Lett.* **1989**, *157* (6), 479–483.
- (11) Bore, S. L.; Paesani, F. Realistic Phase Diagram of Water from “First Principles” Data-Driven Quantum Simulations. *Nat. Commun.* **2023**, *14* (1), 3349.
- (12) Baer, M. D.; Mundy, C. J. Toward an Understanding of the Specific Ion Effect Using Density Functional Theory. *J. Phys. Chem. Lett.* **2011**, *2* (9), 1088–1093.
- (13) Kocer, E.; Ko, T. W.; Behler, J. Neural Network Potentials: A Concise Overview of Methods. *Annu. Rev. Phys. Chem.* **2022**, *73*, 163–186.
- (14) Behler, J. Four Generations of High-Dimensional Neural Network Potentials. *Chem. Rev.* **2021**, *121* (16), 10037–10072.
- (15) Anstine, D.; Zubatyuk, R.; Isayev, O. AIMNet2: A Neural Network Potential to Meet Your Neutral, Charged, Organic, and Elemental-Organic Needs. *ChemRxiv* **2023**, 1 DOI: [10.26434/chemrxiv-2023-296ch](https://doi.org/10.26434/chemrxiv-2023-296ch).
- (16) Zhang, J.; Pagotto, J.; Duignan, T. T. Towards Predictive Design of Electrolyte Solutions by Accelerating *Ab Initio* Simulation with Neural Networks. *J. Mater. Chem. A* **2022**, *10* (37), 19560–19571.
- (17) Baker, S.; Pagotto, J.; Duignan, T. T.; Page, A. J. High-Throughput Aqueous Electrolyte Structure Prediction Using IonSolvR and Equivariant Graph Neural Network Potentials. *J. Phys. Chem. Lett.* **2023**, *14* (42), 9508–9515.
- (18) Sadus, R. J. Molecular Simulation Meets Machine Learning. *J. Chem. Eng. Data* **2024**, *69* (1), 3–11.
- (19) Daru, J.; Forbert, H.; Behler, J.; Marx, D. Coupled Cluster Molecular Dynamics of Condensed Phase Systems Enabled by Machine Learning Potentials: Liquid Water Benchmark. *Phys. Rev. Lett.* **2022**, *129* (22), No. 226001.
- (20) Weiler, M.; Geiger, M.; Welling, M.; Boomsma, W.; Cohen, T. 3D Steerable CNNs: Learning Rotationally Equivariant Features in Volumetric Data. In *Advances in Neural Information Processing Systems*; NeurIPS, 2018; Vol. 31.
- (21) Batzner, S.; Musaelian, A.; Sun, L.; Geiger, M.; Mailoa, J. P.; Kornbluth, M.; Molinari, N.; Smidt, T. E.; Kozinsky, B. E(3)-Equivariant Graph Neural Networks for Data-Efficient and Accurate Interatomic Potentials. *Nat. Commun.* **2022**, *13*, 2453.
- (22) Batzner, S.; Musaelian, A.; Kozinsky, B. Advancing Molecular Simulation with Equivariant Interatomic Potentials. *Nat. Rev. Phys.* **2023**, *5* (8), 437–438.

- (23) Maxson, T.; Szilvási, T. Transferable Water Potentials Using Equivariant Neural Networks. *arXiv* **2024**, 1.
- (24) Bouatta, N.; Sorger, P.; Alquraishi, M. Feature Articles Protein Structure Prediction by AlphaFold 2 : Are Attention and Symmetries All You Need ? *Structural Biology* **2021**, *77*, 982–991.
- (25) Duval, A.; Mathis, S. V.; Joshi, C. K.; Schmidt, V.; Miret, S.; Malliaros, F. D.; Cohen, T.; Lio, P.; Bengio, Y.; Bronstein, M. A Hitchhiker's Guide to Geometric GNNs for 3D Atomic Systems. *arXiv* **2023**, 1.
- (26) Mijsch, A. M.; Morawietz, T.; Kästner, J.; Urban, A.; Artrith, N. Strategies for the Construction of Machine-Learning Potentials for Accurate and Efficient Atomic-Scale Simulations. *Mach. Learn.: Sci. Technol.* **2021**, *2* (3), No. 031001.
- (27) Jacobson, L.; Stevenson, J.; Ramezanghorbani, F.; Dajnowicz, S.; Leswing, K. Leveraging Multitask Learning to Improve the Transferability of Machine Learned Force Fields. *ChemRxiv* **2023**, 1.
- (28) Batatia, I.; Benner, P.; Chiang, Y.; Elena, A. M.; Kovács, D. P.; Riebesell, J.; Advincula, X. R.; Asta, M.; Baldwin, W. J.; Bernstein, N.; Bhowmik, A.; Blau, S. M.; Cărare, V.; Darby, J. P.; De, S.; Della Pia, F.; Deringer, V. L.; Elijošius, R.; El-Machachi, Z.; Fako, E.; Ferrari, A. C.; Genreith-Schriever, A.; George, J.; Goodall, R. E. A.; Grey, C. P.; Han, S.; Handley, W.; Heenen, H. H.; Hermansson, K.; Holm, C.; Jaafar, J.; Hofmann, S.; Jakob, K. S.; Jung, H.; Kapil, V.; Kaplan, A. D.; Karimitari, N.; Kroupa, N.; Kullgren, J.; Kuner, M. C.; Kuryla, D.; Liepuoniute, G.; Margraf, J. T.; Magdau, I.-B.; Michaelides, A.; Moore, J. H.; Naik, A. A.; Niblett, S. P.; Norwood, S. W.; O'Neill, N.; Ortner, C.; Persson, K. A.; Reuter, K.; Rosen, A. S.; Schaaf, L. L.; Schran, C.; Sivonxay, E.; Stenzel, T. K.; Svahn, V.; Sutton, C.; van der Oord, C.; Varga-Umbrich, E.; Vegge, T.; Vondrák, M.; Wang, Y.; Witt, W. C.; Zills, F.; Csányi, G. A Foundation Model for Atomistic Materials Chemistry. *arXiv* **2023**, 1.
- (29) Merchant, A.; Batzner, S.; Schoenholz, S. S.; Aykol, M.; Cheon, G.; Cubuk, E. D. Scaling Deep Learning for Materials Discovery. *Nature* **2023**, *624*, 80–85.
- (30) Kolluru, A.; Shuaibi, M.; Palizhati, A.; Shoghi, N.; Das, A.; Wood, B.; Zitnick, C. L.; Kitchin, J. R.; Ulissi, Z. W. Open Challenges in Developing Generalizable Large-Scale Machine-Learning Models for Catalyst Discovery. *ACS Catal.* **2022**, *12* (14), 8572–8581.
- (31) Eyert, V.; Wormald, J.; Curtin, W. A.; Wimmer, E. Machine-Learned Interatomic Potentials: Recent Developments and Prospective Applications. *J. Mater. Res.* **2023**, *38* (24), 5079–5094.
- (32) Omranpour, A.; De Hijes, P. M.; Behler, J.; Dellago, C. Perspective: Atomistic Simulations of Water and Aqueous Systems with Machine Learning Potentials. *arXiv* **2024**, 1.
- (33) Yang, M.; Raucci, U.; Parrinello, M. Reactant-Induced Dynamics of Lithium Imide Surfaces during the Ammonia Decomposition Process. *Nat. Catal.* **2023**, *6* (9), 829–836.
- (34) Zariquiey, F. S.; Galvelis, R.; Gallicchio, E.; Chodera, J. D.; Markland, T. E.; de Fabritiis, G. Enhancing Protein-Ligand Binding Affinity Predictions Using Neural Network Potentials. *arXiv* **2024**, 1.
- (35) Song, S.; Vuckovic, S.; Kim, Y.; Yu, H.; Sim, E.; Burke, K. Extending Density Functional Theory with near Chemical Accuracy beyond Pure Water. *Nat. Commun.* **2023**, *14*, 799.
- (36) Palos, E.; Caruso, A.; Paesani, F. Consistent Density Functional Theory-Based Description of Ion Hydration Through Density-Corrected Many-Body Representations. *ChemRxiv* **2023**, 1.
- (37) Palos, E.; Lambros, E.; Sweet, S.; Hu, J.; Dasgupta, S.; Paesani, F. Assessing the Interplay between Functional-Driven and Density-Driven Errors in DFT Models of Water. *J. Chem. Theory Comput.* **2022**, *18* (6), 3410–3426.
- (38) Dasgupta, S.; Shahi, C.; Bhetwal, P.; Perdew, J. P.; Paesani, F. How Good Is the Density-Corrected SCAN Functional for Neutral and Ionic Aqueous Systems, and What Is so Right about the Hartree-Fock Density ? *J. Chem. Theory Comput.* **2022**, *18* (8), 4745–4761.
- (39) Belleflamme, F.; Hutter, J. Radicals in Aqueous Solution: Assessment of Density-Corrected SCAN Functional. *Phys. Chem. Chem. Phys.* **2023**, *25* (31), 20817–20836.
- (40) Pagotto, J.; Zhang, J.; Duignan, T. T. Predicting the Properties of Salt Water Using Neural Network Potentials and Continuum Solvent Theory. *ChemRxiv* **2022**, 1 DOI: 10.26434/chemrxiv-2022-jndlx.
- (41) Zhang, J.; Pagotto, J.; Gould, T.; Duignan, T. T. Accurate, Fast and Generalisable First Principles Simulation of Aqueous Lithium Chloride. *arXiv* **2023**, 1 DOI: 10.48550/arXiv.2310.12535.
- (42) Yue, S.; Muniz, M. C.; Calegari Andrade, M. F.; Zhang, L.; Car, R.; Panagiotopoulos, A. Z. When Do Short-Range Atomistic Machine-Learning Models Fall Short? *J. Chem. Phys.* **2021**, *154*, No. 034111.
- (43) Zhai, Y.; Caruso, A.; Bore, S. L.; Luo, Z.; Paesani, F. A “Short Blanket” Dilemma for a State-of-the-Art Neural Network Potential for Water: Reproducing Experimental Properties or the Physics of the Underlying Many-Body Interactions? *J. Chem. Phys.* **2023**, *158*, No. 084111.
- (44) Yao, K.; Herr, J. E.; Toth, D. W.; Mckintyre, R.; Parkhill, J. The TensorMol-0.1 Model Chemistry: A Neural Network Augmented with Long-Range Physics. *Chem. Sci.* **2018**, *9* (8), 2261–2269.
- (45) Kosmala, A.; Gasteiger, J.; Gao, N.; Günnemann, S. Ewald-Based Long-Range Message Passing for Molecular Graphs. *arXiv* **2023**, 1 DOI: 10.48550/arXiv.2303.04791.
- (46) Anstine, D. M.; Isayev, O. Machine Learning Interatomic Potentials and Long-Range Physics. *J. Phys. Chem. A* **2023**, *127*, No. 2417.
- (47) Yu, H.; Hong, L.; Chen, S.; Gong, X.; Xiang, H. Capturing Long-Range Interaction with Reciprocal Space Neural Network. *arXiv* **2022**, 1 DOI: 10.48550/arXiv.2211.16684.
- (48) Grisafi, A.; Ceriotti, M. Incorporating Long-Range Physics in Atomic-Scale Machine Learning. *J. Chem. Phys.* **2019**, *151*, 204105.
- (49) Gao, A.; Remsing, R. C. Self-Consistent Determination of Long-Range Electrostatics in Neural Network Potentials. *Nat. Commun.* **2022**, *13*, 1572.
- (50) Zhang, L.; Wang, H.; Muniz, M. C.; Panagiotopoulos, A. Z.; Car, R.; E, W. A Deep Potential Model with Long-Range Electrostatic Interactions. *J. Chem. Phys.* **2022**, *156*, No. 124107.
- (51) Ko, T. W.; Finkler, J. A.; Goedecker, S.; Behler, J. Accurate Fourth-Generation Machine Learning Potentials by Electrostatic Embedding. *arXiv* **2023**, 1.
- (52) Niblett, S. P.; Galib, M.; Limmer, D. T. Learning Intermolecular Forces at Liquid-Vapor Interfaces. *J. Chem. Phys.* **2021**, *155*, No. 164101.
- (53) Zhang, L.; Lin, D.-Y.; Wang, H.; Car, R.; E, W. Active Learning of Uniformly Accurate Interatomic Potentials for Materials Simulation. *Physical Review Materials* **2019**, *3*, No. 023804.
- (54) Guo, J.; Woo, V.; Andersson, D.; Hoyt, N.; Williamson, M.; Foster, I.; Benmore, C.; Jackson, N.; Sivaraman, G. AL4GAP: Active Learning Workflow for Generating DFT-SCAN Accurate Machine-Learning Potentials for Combinatorial Molten Salt Mixtures. *ChemRxiv* **2023**, 1.
- (55) Kulichenko, M.; Barros, K.; Lubbers, N.; Li, Y. W.; Messerly, R.; Tretiak, S.; Smith, J. S.; Nebgen, B. Uncertainty-Driven Dynamics for Active Learning of Interatomic Potentials. *Nat. Comput. Sci.* **2023**, *3* (3), 230–239.
- (56) Gregory, K. P.; Elliott, G. R.; Wanless, E. J.; Webber, G. B.; Page, A. J. A Quantum Chemical Molecular Dynamics Repository of Solvated Ions. *Scientific Data* **2022**, *9*, 430.
- (57) Guo, Z.; Lu, D.; Yan, Y.; Hu, S.; Liu, R.; Tan, G.; Sun, N.; Jiang, W.; Liu, L.; Chen, Y.; Zhang, L.; Chen, M.; Wang, H.; Jia, W. Extending the Limit of Molecular Dynamics with Ab Initio Accuracy to 10 Billion Atoms. *arXiv* **2022**, 1.
- (58) Musaelian, A.; Johansson, A.; Batzner, S.; Kozinsky, B. Scaling the Leading Accuracy of Deep Equivariant Models to Biomolecular Simulations of Realistic Size. *arXiv* **2023**, 1.
- (59) Nguyen, T. T.; Székely, E.; Imbalzano, G.; Behler, J.; Csányi, G.; Ceriotti, M.; Götz, A. W.; Paesani, F. Comparison of Permutationally Invariant Polynomials, Neural Networks, and Gaussian Approximation Potentials in Representing Water Interactions through Many-Body Expansions. *J. Chem. Phys.* **2018**, *148*, 241725.
- (60) de Hijes, P. M.; Dellago, C.; Jinnouchi, R.; Schmiedmayer, B.; Kresse, G. Comparing Machine Learning Potentials for Water: Kernel-Based Regression and Behler-Parrinello Neural Networks. *arXiv* **2023**, 1.

- (61) Pozdnyakov, S. N.; Ceriotti, M. Smooth, Exact Rotational Symmetrization for Deep Learning on Point Clouds. *arXiv* **2023**, 1.
- (62) Takaba, K.; Pulido, L.; Behara, P. K.; Cavender, C. E.; Friedman, A. J.; Henry, M. M.; MacDermott-Opeskin, H.; Iacovella, C. R.; Nagle, A. M.; Payne, A. M.; et al. Espaloma-0.3.0: Machine-Learned Molecular Mechanics Force Field for the Simulation of Protein-Ligand Systems and Beyond. *arXiv* **2023**, 1.
- (63) Musaelian, A.; Batzner, S.; Johansson, A.; Sun, L.; Owen, C. J.; Kornbluth, M.; Kozinsky, B. Learning Local Equivariant Representations for Large-Scale Atomistic Dynamics. *Nat. Commun.* **2023**, *14*, 579.
- (64) Pelaez, R. P.; Simeon, G.; Galvelis, R.; Mirarchi, A.; Eastman, P.; Doerr, S.; Thölke, P.; Markland, T. E.; De Fabritiis, G. TorchMD-Net 2.0: Fast Neural Network Potentials for Molecular Simulations. *arXiv* **2024**, 1.
- (65) Wang, J.; Olsson, S.; Wehmeyer, C.; Pérez, A.; Charron, N. E.; De Fabritiis, G.; Noé, F.; Clementi, C. Machine Learning of Coarse-Grained Molecular Dynamics Force Fields. *ACS Cent. Sci.* **2019**, *5* (5), 755–767.
- (66) Chen, Y.; Krämer, A.; Charron, N. E.; Husic, B. E.; Clementi, C.; Noé, F. Machine Learning Implicit Solvation for Molecular Dynamics. *J. Chem. Phys.* **2021**, *155*, No. 084101.
- (67) Durumeric, A. E. P.; Charron, N. E.; Templeton, C.; Musil, F.; Bonneau, K.; Pasos-Trejo, A. S.; Chen, Y.; Kelkar, A.; Noé, F.; Clementi, C. Machine Learned Coarse-Grained Protein Force-Fields: Are We There Yet? *Curr. Opin. Struct. Biol.* **2023**, *79*, No. 102533.
- (68) Wilson, M. O.; Huang, D. M. Anisotropic Molecular Coarse-Graining by Force and Torque Matching with Neural Networks. *J. Chem. Phys.* **2023**, *159*, No. 024110.
- (69) Wellawatte, G. P.; Hocky, G. M.; White, A. D. Neural Potentials of Proteins Extrapolate beyond Training Data. *J. Chem. Phys.* **2023**, *159*, No. 085103.
- (70) Majewski, M.; Pérez, A.; Thölke, P.; Doerr, S.; Charron, N. E.; Giorgino, T.; Husic, B. E.; Clementi, C.; Noé, F.; De Fabritiis, G. Machine Learning Coarse-Grained Potentials of Protein Thermodynamics. *Nat. Commun.* **2023**, *14*, 5739.
- (71) Loose, T. D.; Sahrman, P. G.; Qu, T. S.; Voth, G. A. Coarse-Graining with Equivariant Neural Networks: A Path Toward Accurate and Data-Efficient Models. *J. Phys. Chem. B* **2023**, *127* (49), 10564–10572.
- (72) Loose, T. D.; Sahrman, P. G.; Qu, T. S.; Voth, G. A. Coarse-Graining with Equivariant Neural Networks: A Path Toward Accurate and Data-Efficient Models. *J. Phys. Chem. B* **2023**, *127*, No. 10564.
- (73) Charron, N. E.; Musil, F.; Guljas, A.; Chen, Y.; Bonneau, K.; Pasos-Trejo, A. S.; Venturin, J.; Gusew, D.; Zaporozhets, I.; Krämer, A.; Templeton, C.; Kelkar, A.; Durumeric, A. E. P.; Olsson, S.; Pérez, A.; Majewski, M.; Husic, B. E.; Patel, A.; De Fabritiis, G.; Noé, F.; Clementi, C. Navigating Protein Landscapes with a Machine-Learned Transferable Coarse-Grained Model. *arXiv* **2023**, 1.
- (74) Li, Z.; Wellawatte, G. P.; Chakraborty, M.; Gandhi, H. A.; Xu, C.; White, A. D. Graph Neural Network Based Coarse-Grained Mapping Prediction. *Chemical Science* **2020**, *11* (35), 9524–9531.
- (75) Wang, W.; Xu, M.; Cai, C.; Miller, B. K.; Smidt, T.; Wang, Y.; Tang, J.; Gómez-Bombarelli, R. Generative Coarse-Graining of Molecular Conformations. *arXiv* **2022**, 1.
- (76) Jin, J.; Schweizer, K. S.; Voth, G. A. Understanding Dynamics in Coarse-Grained Models. I. Universal Excess Entropy Scaling Relationship. *J. Chem. Phys.* **2023**, *158*, No. 034103.
- (77) Dalton, B. A.; Ayaz, C.; Kiefer, H.; Klimek, A.; Tepper, L.; Netz, R. R. Fast Protein Folding Is Governed by Memory-Dependent Friction. *Proc. Natl. Acad. Sci. U.S.A.* **2023**, *120* (31), No. e2220068120.
- (78) Gould, T.; Kronik, L. Ensemble Generalized Kohn–Sham Theory: The Good, the Bad, and the Ugly. *J. Chem. Phys.* **2021**, *154*, No. 094125.
- (79) Riplinger, C.; Neese, F. An Efficient and near Linear Scaling Pair Natural Orbital Based Local Coupled Cluster Method. *J. Chem. Phys.* **2013**, *138*, No. 034106.
- (80) Stein, F.; Hutter, J. Massively Parallel Implementation of Gradients within the Random Phase Approximation: Application to the Polymorphs of Benzene. *J. Chem. Phys.* **2024**, *160*, No. 024120.
- (81) White, S. R.; Martin, R. L. *Ab Initio* Quantum Chemistry Using the Density Matrix Renormalization Group. *J. Chem. Phys.* **1999**, *110* (9), 4127–4130.
- (82) Hermann, J.; Spencer, J.; Choo, K.; Mezzacapo, A.; Foulkes, W. M. C.; Pfau, D.; Carleo, G.; Noé, F. Ab-Initio Quantum Chemistry with Neural-Network Wavefunctions. *arXiv* **2022**, 1.
- (83) Choo, K.; Neupert, T.; Carleo, G. Two-Dimensional Frustrated J1 – J2 Model Studied with Neural Network Quantum States. *Phys. Rev. B* **2019**, *100* (12), No. 125124.
- (84) Carrasquilla, J.; Torlai, G. How To Use Neural Networks To Investigate Quantum Many-Body Physics. *PRX Quantum* **2021**, *2* (4), No. 040201.
- (85) M Casares, P. A.; Baker, J. S.; Medvidović, M.; Reis, R. D.; Arrazola, J. M. GradDFT. A Software Library for Machine Learning Enhanced Density Functional Theory. *J. Chem. Phys.* **2024**, *160*, No. 062501.
- (86) Lee, S.; Lee, J.; Zhai, H.; Tong, Y.; Dalzell, A. M.; Kumar, A.; Helms, P.; Gray, J.; Cui, Z.-H.; Liu, W.; Kastoryano, M.; Babbush, R.; Preskill, J.; Reichman, D. R.; Campbell, E. T.; Valeev, E. F.; Lin, L.; Chan, G. K.-L. Is There Evidence for Exponential Quantum Advantage in Quantum Chemistry? *arXiv* **2022**, 1.
- (87) Waintal, X. The Quantum House Of Cards. *Proc. Natl. Acad. Sci. U.S.A.* **2024**, *121* (1), No. e2313269120.
- (88) Arts, M.; Garcia Satorras, V.; Huang, C.-W.; Zügner, D.; Federici, M.; Clementi, C.; Noé, F.; Pinsler, R.; Van Den Berg, R. Two for One: Diffusion Models and Force Fields for Coarse-Grained Molecular Dynamics. *J. Chem. Theory Comput.* **2023**, *19* (18), 6151–6159.
- (89) Watson, J. L.; Juergens, D.; Bennett, N. R.; Trippe, B. L.; Yim, J.; Eisenach, H. E.; Ahern, W.; Borst, A. J.; Ragotte, R. J.; Milles, L. F.; Wicky, B. I. M.; Hanikel, N.; Pellock, S. J.; Courbet, A.; Sheffler, W.; Wang, J.; Venkatesh, P.; Sappington, I.; Torres, S. V.; Lauko, A.; De Bortoli, V.; Mathieu, E.; Ovchinnikov, S.; Barzilay, R.; Jaakkola, T. S.; DiMaio, F.; Baek, M.; Baker, D. De Novo Design of Protein Structure and Function with RFDiffusion. *Nature* **2023**, *620* (7976), 1089–1100.
- (90) Zeni, C.; Pinsler, R.; Zügner, D.; Fowler, A.; Horton, M.; Fu, X.; Shysheya, S.; Crabbé, J.; Sun, L.; Smith, J.; Tomioka, R.; Xie, T. MatterGen: A Generative Model for Inorganic Materials Design. *arXiv* **2023**, 1.
- (91) Fowler, N. J.; Williamson, M. P. The Accuracy of Protein Structures in Solution Determined by AlphaFold and NMR. *Structure* **2022**, *30* (7), 925.
- (92) Burley, S. K.; Berman, H. M.; Bhikadiya, C.; Bi, C.; Chen, L.; Costanzo, L. D.; Christie, C.; Duarte, J. M.; Dutta, S.; Feng, Z.; Ghosh, S.; Goodsell, D. S.; Green, R. K.; Guranovic, V.; Guzenko, D.; Hudson, B. P.; Liang, Y.; Lowe, R.; Peisach, E.; Periskova, I.; Randle, C.; Rose, A.; Sekharan, M.; Shao, C.; Tao, Y.-P.; Valasatava, Y.; Voigt, M.; Westbrook, J.; Young, J.; Zardecki, C.; Zhuravleva, M.; Kurisu, G.; Nakamura, H.; Kengaku, Y.; Cho, H.; Sato, J.; Kim, J. Y.; Ikegawa, Y.; Nakagawa, A.; Yamashita, R.; Kudou, T.; Bekker, G.-J.; Suzuki, H.; Iwata, T.; Yokochi, M.; Kobayashi, N.; Fujiwara, T.; Velankar, S.; Kleywegt, G. J.; Anyango, S.; Armstrong, D. R.; Berrisford, J. M.; Conroy, M. J.; Dana, J. M.; Deshpande, M.; Gane, P.; Gáborová, R.; Gupta, D.; Gutmanas, A.; Koča, J.; Mak, L.; Mir, S.; Mukhopadhyay, A.; Nadzirin, N.; Nair, S.; Patwardhan, A.; Paysan-Lafosse, T.; Pravda, L.; Salih, O.; Sehnal, D.; Varadi, M.; Váreková, R.; Markley, J. L.; Hoch, J. C.; Romero, P. R.; Baskaran, K.; Maziuk, D.; Ulrich, E. L.; Wedell, J. R.; Yao, H.; Livny, M.; Ioannidis, Y. E. Protein Data Bank: The Single Global Archive for 3D Macromolecular Structure Data. *Nucleic Acids Res.* **2019**, *47* (D1), D520–D528.
- (93) Chakravarty, D.; Schafer, J. W.; Chen, E. A.; Thole, J. R.; Porter, L. L. AlphaFold2 Has More to Learn about Protein Energy Landscapes. *BioRxiv* **2023**, 1.
- (94) Roney, J. P.; Ovchinnikov, S. State-of-the-Art Estimation of Protein Model Accuracy Using AlphaFold. *Phys. Rev. Lett.* **2022**, *129* (23), No. 238101.
- (95) Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; Lillicrap, T.; Simonyan, K.; Hassabis, D. A General Reinforcement Learning

Algorithm That Masters Chess, Shogi, and Go through Self-Play. *Science* **2018**, 362 (6419), 1140–1144.

(96) Trinh, T. H.; Wu, Y.; Le, Q. V.; He, H.; Luong, T. Solving Olympiad Geometry without Human Demonstrations. *Nature* **2024**, 625 (7995), 476–482.

(97) Reiser, P.; Neubert, M.; Eberhard, A.; Torresi, L.; Zhou, C.; Shao, C.; Metni, H.; van Hoesel, C.; Schopmans, H.; Sommer, T.; Friederich, P. Graph Neural Networks for Materials Science and Chemistry. *Commun. Mater.* **2022**, 3, 93.

(98) Wang, W.; Wu, Z.; Dietschreit, J. C. B.; Gómez-Bombarelli, R. Learning Pair Potentials Using Differentiable Simulations. *J. Chem. Phys.* **2023**, 158, No. 044113.

(99) Kaufman, B.; Williams, E. C.; Underkoffler, C.; Pederson, R.; Mardirossian, N.; Watson, I.; Parkhill, J. COATI: Multimodal Contrastive Pretraining for Representing and Traversing Chemical Space. *J. Chem. Inf. Model.* **2024**, 64 (4), 1145–1157.

(100) White, A. D. The Future of Chemistry Is Language. *Nature Reviews Chemistry* **2023**, 7, 457–458.