

# What the foundations of quantum computer science teach us about chemistry

Cite as: J. Chem. Phys. 155, 150901 (2021); <https://doi.org/10.1063/5.0060367>

Submitted: 17 June 2021 • Accepted: 20 September 2021 • Published Online: 15 October 2021

 Jarrod R. McClean, Nicholas C. Rubin,  Joonho Lee, et al.



View Online




Export Citation



CrossMark

Challenge us.

What are your needs for  
periodic signal detection?



Zurich  
Instruments

# What the foundations of quantum computer science teach us about chemistry

Cite as: J. Chem. Phys. 155, 150901 (2021); doi: 10.1063/5.0060367

Submitted: 17 June 2021 • Accepted: 20 September 2021 •

Published Online: 15 October 2021






View Online



Export Citation



CrossMark

Jarrold R. McClean,<sup>1,a)</sup>  Nicholas C. Rubin,<sup>1</sup> Joonho Lee,<sup>1,2</sup>  Matthew P. Harrigan,<sup>1</sup>  Thomas E. O'Brien,<sup>1</sup> Ryan Babbush,<sup>1</sup> William J. Huggins,<sup>1</sup> and Hsin-Yuan Huang<sup>3</sup>

## AFFILIATIONS

<sup>1</sup>Google Quantum AI, 340 Main Street, Venice, California 90291, USA

<sup>2</sup>Department of Chemistry, Columbia University, New York, New York 10027, USA

<sup>3</sup>Institute for Quantum Information and Matter and Department of Computing and Mathematical Sciences, Caltech, Pasadena, California 91125, USA

<sup>a)</sup>Author to whom correspondence should be addressed: [jmcclean@google.com](mailto:jmcclean@google.com)

## ABSTRACT

With the rapid development of quantum technology, one of the leading applications that has been identified is the simulation of chemistry. Interestingly, even before full scale quantum computers are available, quantum computer science has exhibited a remarkable string of results that directly impact what is possible in a chemical simulation with any computer. Some of these results even impact our understanding of chemistry in the real world. In this Perspective, we take the position that direct chemical simulation is best understood as a digital experiment. While on the one hand, this clarifies the power of quantum computers to extend our reach, it also shows us the limitations of taking such an approach too directly. Leveraging results that quantum computers cannot outpace the physical world, we build to the controversial stance that some chemical problems are best viewed as problems for which no algorithm can deliver their solution, in general, known in computer science as undecidable problems. This has implications for the predictive power of thermodynamic models and topics such as the ergodic hypothesis. However, we argue that this Perspective is not defeatist but rather helps shed light on the success of existing chemical models such as transition state theory, molecular orbital theory, and thermodynamics as models that benefit from data. We contextualize recent results, showing that data-augmented models are a more powerful rote simulation. These results help us appreciate the success of traditional chemical theory and anticipate new models learned from experimental data. Not only can quantum computers provide data for such models, but they can also extend the class and power of models that utilize data in fundamental ways. These discussions culminate in speculation on new ways for quantum computing and chemistry to interact and our perspective on the eventual roles of quantum computers in the future of chemistry.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0060367>

## I. INTRODUCTION

The study of quantum computing in the abstract is an opportunity to ask ourselves what is possible if we could attain an almost unimaginable level of control of the microscopic facets of our universe. It was first proposed as a solution to the problem of simulating physical systems with strongly quantum characteristics, a task that has proven very challenging for traditional computers.<sup>1</sup> The idea was that if, like the puppet of a marionette puppeteer, one could make a precisely controllable quantum system act enough like a more interesting system, then the puppet could answer previously unknown questions about the true system. This core concept of quantum

simulation eventually merged with modern computer science to form the fields of quantum computer science and quantum computing.<sup>2</sup> This merger allowed these concepts to be made more precise and expanded applications beyond physical systems into abstract ones such as breaking cryptography.<sup>3,4</sup>

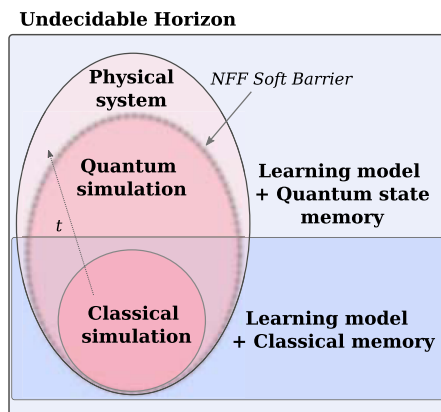
Despite the expansion into other applications, the simulation of quantum systems and, especially, strongly correlated chemistry, has remained a primary application of interest. Chemistry represents a sweet spot of quantum effects strong enough to make them challenging for classical computers while still having a well-known application space to motivate development. Since the original proposal of Aspuru-Guzik *et al.*,<sup>5</sup> there have been a number of developments in

quantum algorithms for the direct simulation of chemistry bringing costs down from astronomical numbers to routine calculations for even very challenging systems.<sup>6–10</sup>

In parallel to algorithmic developments, quantum technology has advanced at a rapid pace. Recent demonstrations by the Google group have experimentally shown that these devices are capable of tasks that are incredibly challenging on a classical computer, with a gap that will only grow with the quality of modern quantum computers.<sup>11</sup> In addition, a number of prototype chemistry experiments have now been experimentally demonstrated on quantum devices.<sup>10,12–14</sup> At present, error rates in quantum devices are too great to make them competitive with the best classical algorithms for chemistry, but quantum error correction and mitigation offer paths forward.<sup>15–19</sup>

However, even without an actual quantum computer running an algorithm, the study of quantum computer science in the abstract has led to insights about the limitations and possibilities of chemical simulation, and, indeed, under modest assumptions, about the universe and chemistry itself. Recent results tying quantum computer science directly to chemistry include showing that the electronic structure problem is expected to be hard in the worst case even for quantum computers.<sup>20–22</sup> In this work, we highlight some of the yet unexplored connections leading up to a perspective on the role both traditional and quantum computers may play in the future of chemistry. To build to this perspective, we begin by framing direct chemical simulation as a digital experiment. In this framework, we exploit known results that restrict the power of any computer, even a quantum one, to understand limits of such digital experiments in chemistry. Along the way, we encounter surprising results about the impossibility of an algorithm for determining if a system ever thermalizes, flying in the face of conventional thermodynamic analysis. However, this construction guides us to new ground, where we will see that learning from data can be fundamentally more powerful than traditional computation. This sheds new light on the success of many existing chemical theories and the way in which they exceed rote simulation by leveraging the additional power provided by data.

The relationship between traditional simulation and learning models in chemistry is cartooned in Fig. 1, where classical models with data from nature can exceed the physical prediction horizons of rote simulation. We consider learning models to be models where some amount of high quality training data for the task is available, from either a physical experiment or a simulation in addition to just the specification of the problem. As we will argue, this is not just a reference to modern machine learning methods but rather encompasses the foundations of chemical theory such as transition state theory, molecular orbital theory, and thermodynamically controlled reactions. In addition, we take the opportunity to highlight previously unused technology from quantum information science that may bolster the development of chemistry even before the arrival of full fault tolerant quantum computers. The links from chemistry to results in quantum computer science all build to our ultimate perspective that the eventual role of quantum computers in chemistry will be to aid in the construction of learning models for chemistry. This includes providing data to classical models, constructing quantum models that can make accurate predictions with far less data, and eventually interfacing directly with quantum data from chemical systems. To close, we wrap this



**FIG. 1.** Cartoon of the relative power of direct simulation vs learning models for chemistry, where a point on the figure is a question related to a chemical system. The ovals on the left depict questions efficiently accessible via simulation of time dynamics, where quantum simulation is widely believed to be exponentially stronger than classical simulation but still softly bound (non-exponential) by the no-fast-forwarding (NFF) to lag behind the system in nature even with a quantum computer. Learning models, which include both traditional theories supported by data such as thermodynamics and modern machine learned models, strictly include all questions answerable by traditional simulation, as they are assumed to not only have access to simulation but also have their power enhanced by data that comes from either nature or other quantum simulations. The availability of quantum states held in memory and full quantum computation extends the questions that may be efficiently answerable beyond that of traditional learning models.

Perspective into an outlook for the interplay between these two exciting areas.

## II. DIGITAL CHEMISTRY EXPERIMENTS WITH QUANTUM COMPUTERS

The idea to draw a distinction between theory and computation or simulation is not new, with suggestions of considering computation as the “third pillar of science,” alongside experiment and theory.<sup>23</sup> In this framework, it can sometimes be more accurate to view simulations as closer to experiments than to theory. Quantum computing, and especially quantum simulation of physical systems, helps to make this even more clear by constructing simulations of the physical world at the quantum level with efficiently refineable (and bounded) accuracy. In addition, limitations on precision for quantum computers<sup>24</sup> to learn physical properties sometimes place them even closer to experiments than classical digital simulations of the same systems. While a discussion on the precise distinctions between simulation and theory deserves more exposition than we can provide here with many gray areas of overlap, we will attempt to provide a working distinction. In particular, simulation will refer to a rote input and output relationship, where given some computational problem and inputs, a fixed output is reached. In contrast, theory will refer to the broader activity of developing and refining models that may go into a simulation, where those refinements could be based on data from physical experiments, simulations, or abstract mathematical constructs. In addition, it is worth distinguishing between data that can be considered as input, such as initial conditions of a dynamical system, and data used to develop a theory

or model, which we will refer to as training data in connection with the growing literature on learning theory.

The most natural setting for a quantum simulation on a quantum computer is that of watching the system step forward in time or quantum time dynamics. Within the field of quantum computing, the term “quantum simulation” is sometimes reserved for time dynamics simulation specifically, and when we refer to rote simulation, it will typically apply to this case. That is, one sets an initial state, a physical system defined by its interactions, and some time to simulate, and the quantum computer performs a computational experiment in mapping the initial state to the final time, where any reasonable observable of the system can be probed. This time dynamics obviously mirrors the most natural settings of the real world. However, despite this natural setting, one will note that a large number of quantum algorithms and theories closely follow that of the electronic structure performed on classical computers, where the focus is on low energy eigenstates of the electronic system,<sup>5,25</sup> rather than exclusively the final state after some fixed evolution. This important deviation already wraps in data from the natural world that tells us in many systems, thermal states are an apt description of a system, and at low temperatures, low energy states are most commonly observed. As we will discuss later, the use of such observational knowledge is strictly more powerful than time dynamics simulation and underpins the great success of many modern chemical theory or simulation methods. A detailed discussion of this point will follow, but for now, we lump both static and dynamic experiments into the category of direct simulation.

By direct simulation, we mean a simulation where the system of interest (often processes such as a reaction) is known, and we seek to use the computational experiment to gain insights into the process that is challenging to access through the experiment, for example, details of a reaction mechanism, transition states, or migration of charge through a system. They represent the bulk of computational experiments, both on classical computers and those proposed on quantum computers, and we draw a distinction between direct simulation and design, often characterized as inverse problems.<sup>26</sup>

For the task of direct simulation of chemical systems, quantum computers have been shown to demonstrate an exponential advantage over their classical counterparts under modest assumptions.<sup>5,27–29</sup> An exponential speedup has the practical implication that some simulations that might have taken longer than the age of the universe could be done in mere seconds. More specifically, full quantum dynamics with no approximations beyond discretization error can be performed on a quantum computer in a time that scales only polynomially in  $t$ ,  $N$ , and  $M$ , where  $t$  is the simulated time,  $M$  is the number of basis functions, and  $N$  is the number of electrons.<sup>9,19,30</sup> Interestingly, along with these speedups, it comes with some of the same limitations possessed by a physical experiment. For example, in contrast to classical simulations where more precision in a simulated quantity is often relatively easily obtained, the Heisenberg limit applies to any measurements one might perform to extract information. These approaches have also included exploring the ansatz for electronic systems that are inaccessible to classical computers.<sup>13,31</sup> Perhaps surprisingly, recent results have shown that quantum computers can even achieve a scaling for exact computations that is sub-linear in the basis set size by taking advantage of a first quantized representation in a way that has no current classical counterpart.<sup>9</sup>

However, quantum computers do not vastly extend their capabilities into the realm of discovery or design, which we identify as distinct from direct simulation. That is, while many of the proposed simulations would be faster and more accurate representations of the same systems than would be achievable classically, they are much the same type of the experiment without reaching into the design space. While there has been some notion of how quantum search may assist in design, most results promise at most a quadratic speedup in contrast to exponential speedups in direct simulation,<sup>32</sup> although we note that the direct simulation subroutines used in a potential search would still benefit from improved speed or accuracy. Due to the immense size of the design space, this quadratic speedup is not terribly compelling if it is not combined with structured strategies. That is, any improved search must start not from naive global search but a search enriched by knowledge of the design space, for example, one that starts in the chemical neighborhood of known, synthesizable compounds. In addition, practical issues arising in real quantum computers may prevent quadratic speedups from being advantageous for some time.<sup>33</sup> We will argue later that perhaps it is even more apt and useful to consider an infinite chemical space as opposed to one which is merely combinatorially large.

That said, direct simulation, of course, has a great value for design. For example, if a key reaction or catalyst is known to work to catalyze a reaction of interest, the microscopic detail of a computational experiment that is inaccessible in the lab can offer insights into the causal mechanism that enable the design of analogs.<sup>34–36</sup> In addition, a now popular approach is the use of screening of many candidate compounds through direct simulation in order to identify new candidate molecules or learn properties, which are predictive for a desired design task.<sup>37–39</sup> This type of approach is also used for selection of candidates in protein to small-molecule docking.<sup>40</sup> Direct simulation can also be used for novel discovery of processes through direct time evolution, as has been done with the molecular nano-reactor.<sup>41</sup> Quantum computers serve to strictly improve this type of discovery by expanding the set of systems that can be accurately treated in a given amount of time by such methods. Advantage in improving the accuracy in direct simulation has been the primary focus in much of the literature of quantum computation for chemistry, with the hope that it enables for an accurate study of important systems inaccessible to current methods. The example receiving perhaps the most attention is FeMoCo,<sup>19,28,30,42–44</sup> where the hope is that direct simulation will offer a mechanistic insight that leads to the design of new catalysts for nitrogen fixation. However, when leveraging direct simulations results for an improved design, the simulations are typically tightly coupled to a theory already rooted in data. For example, the energy of a transition state might be used to justify a reaction rate, but this is inherently coupled to the empirical success of transition state theory, rather than a result that emerges directly from simulation.

In contrast, comparatively little attention has been paid to the potential advantages of improved dynamics simulations, perhaps due to the larger foundation of work available utilizing static eigenstates in the classical case. In addition, it is perhaps surprising that dynamics simulations on quantum computers that fully relax the Born–Oppenheimer approximation, using fully quantum nuclei and exact electronic representations, are about as costly as those that live within such approximations. However, with each step toward this nuclear-electron soup, what is gained in accuracy might be paid

in ease of interpretation as the added overhead of re-discovering the notion of a molecule through observable measurements remains a largely uncharted problem in the context of quantum computing.<sup>27,45</sup> Taking advantage of these departures from the tried and true anchors of chemical space and embracing the idea of a fully *ab initio* experiment will require more development in how this type of simulation can be most beneficial. However, if it is unlocked, it may offer new insights into complex dynamic processes in spectroscopic or electrochemical processes where both excited electronic and quantum nuclear states can play a role.<sup>41,46</sup> It may also offer new insights into the long studied challenges surrounding proton coupled electron transfer.<sup>47</sup>

Given the immense power of quantum computers in speeding up these types of direct simulation, one is often left wondering how far these advantages extend. This question is one that has been studied in great detail in computer science, and we now wish to turn to our understanding of the limits of power in quantum computation. The similarity between quantum computations and the physical world implies that the results also tell us something about nature and chemistry itself, and we now turn to this Perspective.

### III. LIMITS OF EVEN QUANTUM COMPUTATIONS IN CHEMISTRY

Having set up the perspective of quantum computation as digital experiments, we now turn to the implications this may have for the physical systems themselves. Due, in part, to the challenge of building scalable quantum hardware, for a good part of its history, quantum computing was a largely theoretical endeavor that aimed at understanding the power of quantum devices and, by extension, physical systems. Considerable progress has been made on this problem, and we wish to highlight those results we believe have the most bearing on chemistry.

Before discussing how they pertain to chemistry, it is worth mentioning a few aspects of the assumptions made in stating the results below. In particular, in order to talk about scaling and cost, it is necessary to have some model of computation where cost can be quantified. We will generally be assuming that this means we use a finite, often local, set of operations called quantum gates, and we work within a digital gate model of quantum computation. This assumption is in some ways analogous to assuming that nature favors local interactions or that computers are built from modular components. Moreover, some of the results below most precisely pertain to the number of times some information, such as the Hamiltonian of a system, must be accessed. This is known as the query complexity, and while it is often closely related to time complexity, these can be different in some circumstances, such as when parallelization is or is not possible and we attempt to be clear as to when this distinction is important.

#### A. Quantum computers cannot fast-forward time either

Starting with the most direct form of quantum simulation, time dynamics of a quantum system, we have already noted the exponential advantage over classical simulation achieved by current quantum algorithms. This begs the question how far these quantum algorithms can go in accelerating simulation of physical systems. For example, can a quantum computer simulate a system faster than nature itself evolves the same system, that is, sub-linear in the time

being simulated? It turns out that, assuming one only has the available space to represent the system serially, the answer to this is a resounding no. It is also worth noting that if a system is finite and the time to simulate is larger than the number of possible system states, the cost of diagonalization, finite state enumeration, or some equivalent can be paid to fast-forward a system. However, in most cases, the time scale of interest is much less than this, and indeed, as we argue later, infinite perspectives on chemical systems are sometimes more natural.

While there are some special systems that can be fast-forwarded,<sup>48</sup> it cannot be done in the most general case. This result can be understood at a high level by imagining if this were true and realizing one could build a recursive simulation that does any computation in vanishingly small time, by nesting simulations within simulations with no space overhead. If one allows the space on the quantum computer to expand polynomially with the simulation time, this result is more unclear and related to the classical debate of if any computation can be parallelized arbitrarily well, encapsulated by the theoretical question of if  $P = NC$ .<sup>49</sup> This is an open question that currently holds a similar status to the more famous question of if  $P = NP$ . Some progress has been made on the question if all computations can be parallelized arbitrarily for specifically quantum computers, showing that more simulation time is strictly more powerful relative to an oracle.<sup>50</sup> To simplify the message here, we assume that quantum space resources are the minimum required to perform the simulation and that as a result, oracle queries must be done in serial. More formally, the limitations of simulation are encapsulated in the no-fast-forwarding theorem.<sup>51,52</sup> This theorem importantly shows that not only is universal computation not fast-forwardable but surprisingly even some Hamiltonians that are incapable of universal computation cannot be fast-forwarded either.

Under the stated assumptions, the no-fast-forwarding result practically implies that no simulation that retains a complete level of detail at the quantum level can simulate physical systems faster than nature, and likely, there will be some large factor disadvantage between them. This result is in some ways not unexpected from classical simulations of classical chemical models such as molecular dynamics, where for a fully detailed tracking of the model in the most general case, one expects to lag behind the physical time scale of the system unless approximations or reduced models are utilized. The emphasis here is that this limitation provably carries into the case of using a quantum computer as well. Taking this result at a face value along with the knowledge that important chemical reactions such as rusting can take months to occur, using a single molecular realization in a simulation could take comparable times to witness a reaction event, making discovery in this fashion totally impractical. However, most chemists would know that is not the end of the story.

Specifically, methods such as transition state theory based on static calculations of free energies or transition path sampling<sup>53</sup> can allow access to rate constants on time scales much faster than direct physical simulations.<sup>54</sup> In addition, the general class of methods related to rare event sampling are often successful at reaching beyond time scales accessible in direct simulation.<sup>55</sup> Such methods have been successful at dramatically increasing the reach of molecular dynamics simulations, often via the observation that rare events are dominated via relatively few pathways. Hence, while the above theorem is true for precise simulations of the full quantum state, we see from practice that physical theories grounded in observation



can circumvent such wild overheads by eliminating the complexity of the full quantum state and focusing on a reduced set of observables, reduced precision, or approximate models amenable to fast-forwarding. Such reductions in the number of important degrees of freedom have been studied in various contexts to try to understand “what makes science possible.”<sup>56,57</sup> It should also be noted that these results are about the most general system, and specific systems may exhibit physics that allow faster than physical simulation when that data are utilized by validating approximate models. In other cases, particular models are amenable to classical fast-forwarding, such as non-interacting electron models, but they are not general enough to encompass all quantum systems with interactions.

The observation that chemical methods have been so successful in making predictions beyond physical time scales leaves one believing that perhaps physical systems always admit some form of reduction that will allow us to forecast results that we actually care about (as opposed to a full quantum state), well in advance of simply stepping them forward in time. A recent string of results in quantum computer science has proven that, in fact, even in some very simple systems by construction, there must exist some questions that are fundamentally unanswerable ahead of time. This list of problems surprisingly includes whether simple one-dimensional (1D) systems thermalize,<sup>58</sup> which has implications for the predictive power of thermodynamics in general systems. These results are much stronger than simple chaos of dynamic systems and require some exposition, but we believe they are a value lens to view chemical problems.

## B. Questions no algorithm can answer and undecidability

A direct reading of the no-fast-forwarding theorem contrasted with the unassailable predictive success of simplified chemical models raises a number of interesting questions. On the one hand, the no-fast-forwarding theorem tells us that even with an exponential advantage over classical simulation of quantum dynamics, a quantum computer aiming to simulate the dynamics of a full wavefunction to a fixed accuracy is bound by some multiplier of the physical time of the actual process. On the other hand, we see that reduced models, for example, mean-field molecular orbital models, reductions based on evaluating the energy of stationary states such as ground and excited states, or even arrow pushing in Lewis dot structures can sometimes be powerfully predictive for a chemical phenomenon of interest at a relative time cost much less than the time scale of a process of interest. This leads one to ask if there is always a predictive model or algorithm for simplified questions. Surprisingly, results in computer science have shown us that there are, indeed, even relatively simple questions related to physical systems for which it can be proven that no single algorithm can answer for all such systems in finite time. Here, while we will certainly not claim that this should cause one to abandon current methods, we will argue that these results should have implications for how one approaches some chemical problems.

To make our case, we will start with a motivating example before we move to the more general concepts that will allow us to address some of the tensions with this interpretation. Consider the question of “Given unbounded time, does this physical system reach thermal equilibrium?” The basic assumptions of conventional thermodynamic theory, leaning on the ergodic hypothesis,<sup>59–61</sup>

assert that the answer is a resounding “yes” under assumptions that the system is ergodic. Indeed, many standard molecular dynamics simulations exploit both possible directions, either determining free energies through proxies derived from long time simulations<sup>54</sup> or predicting long time behavior through models for free energies. However, recent results have shown that there are 1D translationally invariant systems for which no algorithm that can answer the question of thermalization or even simple questions about energy gaps in finite time.<sup>58,62–64</sup>

To make such a problem even more concrete, let us take the question of a finite gap as an example. The problem statement can be set up briefly as follows. One is provided as input of a finite Hamiltonian for a chunk of a system that will be repeated indefinitely. While any finite chunk of this system could be diagonalized, determining the existence of a gap for that chunk, this proof implies that no finite number of chunks taken together could ever tell you about the status of the gap for the infinite system in the general case. Moreover, any extrapolation on those chunks to a large size limit cannot be guaranteed to converge. The statement is so strong that we can guarantee that not only will no simple algorithm work exist, but there cannot exist such an algorithm of any finite time complexity. This type of problem is known as an undecidable problem, the first example of which was the famous halting problem of Turing.<sup>65</sup> Imagining the setup as a physical experiment, one can characterize a small chunk of material experimentally, determining if it has a gap. Assuming that one finds no gap, the size of that material could then be extended, gradually adding more material, and one could again measure at any juncture if the material has a gap. At any fixed size of the problem, one can easily measure the gap, as the undecidable problem does not imply that it is difficult to see that the system is gapless but rather that, in general, one cannot predict whether it will ever exhibit a gap upon measurement at larger sizes in the most general case. This class of problems is much more challenging than even the perhaps more familiar class of NP-complete problems, as even when provided with a supposed answer, a mortal prover cannot verify if it is correct. This implies that for a number such as an energy or free energy representing a system, either that number cannot be generally predictive of the behavior of that system or the number itself must be uncomputable. Said another way, there will be instances where computing the free energy of a chunk of a system fails to capture the complexity of its real environment, and qualitatively different behavior than predicted by thermodynamics or any finite time computation will be observed. By an uncomputable number, we refer to the idea that if there were such a number that could accurately predict the behavior of any system for arbitrary times, its computation would involve the solution of undecidable problems. Showing that this is true, even for a quantum computer, which can perform classical computation as a special case, underscores the weight of these results.

To draw a connection with recent results in worst-case complexity for chemistry, it has been shown that the electronic structure problem in a finite basis and some related problems are in the quantum analog of NP-complete or QMA-complete.<sup>21,22</sup> Such problems live in a finite state space, and hence, the implication is that at worst an exhaustive examination of all possible states will suffice to determine if something will happen akin to a program halting. Despite these complexity results, practicing chemists have obviously not stopped using simulation methods. It is important to remember that most of these results apply for worst-case constructions, and

indeed, physical systems may come from a special set of all possible systems that makes them more amenable to study. For a system with an infinite number of possible states, total enumeration obviously becomes untenable, and some problems graduate to become undecidable. Similar to the case of QMA-completeness, however, this does not imply that every instance of a problem should become impossibly hard.

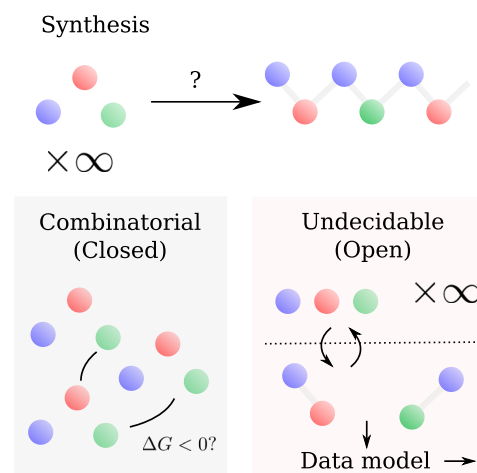
The idea that bare thermodynamic quantities may not be predictive due to lack of thermalization is certainly not new in the study of non-ergodic systems, with preeminent examples such as spontaneous symmetry breaking, spin-glasses, and the general idea of a Gibbs measure.<sup>66-69</sup> These results strengthen these concepts to tell us that in some non-ergodic systems, there is no computable replacement for something like free energy that is predictive of the behavior of a system. Hastily defined, the halting problem is the problem of finding an algorithm that, given a finite program input, can determine in finite time if the program halts or runs forever. If one views a physical evolution as a computation where halting is defined as some qualitative change,<sup>70</sup> for example, the fraction of random systems that thermalize should be closely related to the fraction of random programs that halt, a number known as Chaitin's constant.<sup>71</sup> Indeed, it is easy to see that such a thing requires solving the halting problem for every possible instance to determine the digits of the fraction of halting programs, which defines it as an uncomputable number. For such systems, assuming ergodicity is analogous to solving the halting problem by simply hypothesizing that all programs encoded by the physical evolution halt.

For veteran chemists, another way of viewing this result is that it is a vast strengthening of the concept of kinetically controlled reaction systems to be contrasted with thermodynamically controlled systems. Previous work in this area has shown that this type of unpredictability is distinct and stronger than traditional notion sensitivity induced chaos in rendering systems unpredictable, even when they are deterministic. When considering a computational view of the universe,<sup>2</sup> undecidability does not imply that something cannot happen in the universe. Rather, it implies that in the worst case, some events cannot be predicted using a finite algorithm, even if those events are relatively ordinary and the outcomes are easy to describe once they have occurred. Phrased in a physical point of view, for a system exhibiting undecidable dynamics, there are sudden, qualitative changes that cannot be predicted in any way other than simply stepping the system forward in time.<sup>70</sup> For chemical practitioners, such a description is certainly reminiscent of the emergent phenomena that transition us from simple molecules to proteins, RNA, DNA, complex catalytic networks, and even living systems.

Before continuing, we must address a tension that arises with the mapping of physical processes to halting problems in the form of finite vs infinite systems. In order to map to a halting problem rigorously, one typically maps the operations of a physical system to operations of a Turing machine, with an infinite sized tape. If the tape were finite, one can spend a very long, but not infinite, time examining all possible states of the system to arrive at an answer to any question related to the state of the system. For example, in studying the dynamics of a finite physical system, one could discretize the possible states and expend an exponential amount of time in the system size determining an answer. Hence, one might want to argue that since the universe is often presumed finite, these

conclusions are of little consequence, even though from a practical point of view, the exponential of even a modest sized finite system would take longer than the age of the universe. Despite their somewhat similar predictions, in practice (too long to compute), we argue that embracing the viewpoint of undecidable dynamics allows one to let go of the impossible dream of treating chemical questions by the exponential enumeration of the combinatorial perspective. Whether or not a physical system exhibits truly undecidable behavior, embracing the possibility leads to a more natural description of many chemical processes and points the way toward more apt solutions.

We depict an example of this Perspective in Fig. 2. Consider a collection of small molecules attached to a reservoir of source chemicals, and we wish to know if a particular strand of DNA will ever form, given this setup. If one took the closed system, or combinatorial, viewpoint, it would suggest that perhaps the most effective way to address this question is to model the reservoir with some enormous, but finite, set of explicit molecules and check all the possible states of the dynamic system. Along the way, one can check things such as the change in free energy ( $\Delta G$ ) or simply proceed with short time dynamics. Not only is the approach untenable due to the number of possible configurations, but given the nature of self-catalytic reaction networks, one could accidentally include slightly too few molecules to form the right catalytic process to enable the formation of that strand and reach the wrong conclusion. In contrast, the open system, or undecidable, interpretation of this setup, is that the physical system can continue to borrow from the infinite reservoir such as states of the Turing tape. If it has to temporarily expand to chemical networks of unforeseen size, it can do so. In this situation, one tends



**FIG. 2.** Sketch of two viewpoints for the problem of chemical synthesis. In this scenario, one wants to know if a certain target compound can be reached with a feedstock of specified chemicals under typical conditions. The combinatorial view tries to address this question by introducing a closed reaction system with a large but finite number of the reagents and explores potential reaction pathways to determine feasibility through some proxy such as free energy ( $\Delta G$ ) or a short time simulation. The undecidable viewpoint embraces an open system picture and the relationship to DNA computing to argue that the most feasible approach tries to use data from either nature or related simulations to determine the answer to the question. Such an approach can more naturally include the effects of common side reactions or strong kinetic control.

to lean on data models for what has existed before, and we detail later how doing so changes the power of the simulation approach. The undecidable interpretation is supported by recent work, showing that DNA computing is Turing complete and as such can encode the halting problem.<sup>72,73</sup> While DNA computing is a bit of an artificial construction and real DNA that tends to form in live systems is subject to other constraints that make its activity more predictable, the possibility of implementing universal computation in the dynamics of a chemical system pushes us to embrace the wildest consequences of the theory of computation.

This tension between finite and infinite models is not unfamiliar when related to the study of other non-ergodic systems, such as spin-glasses.<sup>66</sup> In such cases, models might predict that finite systems have some chance of exhibiting substantial rearrangements, albeit on exponentially long time scales. However, it has been seen to be useful to lean on a model that takes a limit where these different basins qualitatively fracture from each other and are considered separate. We suggest here that as molecular systems grow in complexity, the run-away reaction systems that form are best treated by considering them as fracturing into novel parts of chemical space, without tracking vanishingly unlikely events that transition one between these parts of space. Indeed, the undecidability of thermalization suggests that for infinite systems, it is correct to view such fracturing as exact in the most general case. Even in simpler cases, this model tracks more closely with the experience of laboratory chemists who have to contend with side reactions, rendering a thermodynamically favorable synthesis seemingly impossible. By not splitting hairs about the distinction between a large finite system and a truly infinite one, we can move toward practical alternatives to the brute-force combinatorial approach of pure computation.

However, if one resigns themselves to accept that many of these problems have no general algorithm, is one left with no option but to plod forward with rote dynamic simulation or give up? At least the combinatorial approach offered a glimmer of hope, even if it required longer than the age of the universe for tiny problems right? In Sec. IV, we argue that, in fact, this realization does offer an alternative, which focuses the scientist on an approach based on learning and reduced chemical model construction.

#### IV. LEARNING AS AN ALTERNATIVE TO COMPUTATION

Rather than claim that one must give up hope on problems where the most natural formulation looks undecidable or hopelessly expensive, we argue that embracing the undecidable interpretations of these problems frees one from the distraction of the combinatorial approaches and leads them to embracing the only known resolutions to problems that exist for halting problems. The first, and not really palatable solution, that has been mentioned, is running time dynamics forward and hoping for the event of interest to occur. The second, and more actionable viewpoint, is to understand that systems with advice can formally resolve such problems. While advice has a precise definition in theoretical computer science,<sup>74</sup> it was recently shown that data from the real world can act as a restricted form of advice.<sup>75</sup> To solve halting problems, one's advice would need to constitute a form of infinite time pre-computation; however, much like the argument of finite vs infinite systems, the pre-computation performed by the physical universe before this point,

though finite, is already quite strong and available to be revealed through experiments.


To see an example of this in chemistry, one can look to the field of natural product synthesis, where one seeks to synthesize chemical species found in nature from simpler, readily available components in a finite number of steps. For the most general complex molecule, it can be quite difficult to say if a product can be reached with a fixed set of reagents and conditions or at all, perhaps due to some intervening side reaction. In fact, we conjecture this form of the problem may be undecidable since we could specify the target molecule as a particular strand of DNA.<sup>73</sup> By contrast, in the study of natural product synthesis, the challenging question of chemical synthesizability is answered by nature itself, and practitioners are freed to discover practical routes toward a product rather than ponder if success was even possible theoretically. In this sense, empowering ourselves with observations from the physical world can enable answers to questions that cannot be practically resolved with rote algorithms. More generally, we believe one could view the models and study of synthesis in organic chemistry and their predictive success in this light. Hence, taking this point of view suggests that approaches based on learning are fundamentally different from those that rely on computation alone.

With the rise of popularity in machine learning, naturally, there is now a flurry of work applying techniques of machine learning to chemistry.<sup>76</sup> The applications range from prediction of direct simulation quantities<sup>77</sup> to synthesizability<sup>78</sup> and inverse design.<sup>26</sup> We believe that the results we highlight here bolster the motivation for pursuing this line of work in the general sense but also support the long standing tradition of chemists developing physical models as well. While much practical work remains to be done in finding the best representation or approaches, it is now understood that these approaches are fundamentally different, and indeed more powerful in some ways, from traditional simulation approaches.

In particular, recent work has shown that the power offered by data from a quantum computer can lift classical models to be competitive with their quantum counterparts in some cases.<sup>75</sup> If nature is, indeed, a universal quantum computer,<sup>79</sup> this would suggest that quantum computations could help fill the role of natural experiments in providing data to empower learning models, but with improved programmability and flexibility. This might suggest that in the future, a key role of quantum computers, which may still remain scarce compared to their traditional counterparts, may be to provide data for learning models primarily run on traditional computers. However, that would be a rather unexciting fate for a technology that pushes the limits of our understanding of the universe.

Instead of accepting this fate, we appeal to other recent work that has shown an overwhelming advantage for quantum computers in how they can process data moved directly into the computer from a quantum sensor, a process referred to as transduction.<sup>80,81</sup> A comparison with a traditional data pipeline is shown in Fig. 3. In such cases, even performing quantum data processing on even very few copies stored in quantum memory can allow one to extract properties of quantum systems that would require exponentially more data in the classical case.<sup>80</sup> Even early machines may be able to manipulate few copies for limited times and store these quantum data for future models. This separation is stronger than a traditional computational separation, as if limited measurements are available due to the transient nature of a system, no amount of





**FIG. 3.** Comparison of quantum and classical data pipelines. The presence of classical data or traditional measurements, e.g., data from laser spectroscopy or thermodynamic models, already confers an advantage over traditional simulation. Here, we contrast that pipeline with the one where quantum data are collected from quantum sensors, which can be stored in quantum memory, increasing yet again the power of available data. In the classical case, a sensor measures and reports classical values back to a classical computer. In the quantum case, a sensor has its state transferred into a quantum computer, a process also known as transduction, where it may be computed alongside multiple additional copies of such states to perform tasks that would be exponentially more costly on a classical computer. In such setups, the quantum transduction route can require exponentially fewer measurements than the classical case to learn some quantities or perform some tasks. In cases where data are hard to come by, e.g., transient, short-lived, or rare systems, this can make an impossible characterization into the one that is routine.

computation can allow for the traditional measurement and computation to catch up. This implies that even systems with a number of qubits that can be classically simulated could prove advantageous. If the quantum computer has quantum memory that can maintain states in the quantum form, this extends the power of these approaches even further. The relationship between this hierarchy of data models and traditional simulation is depicted in Fig. 1. Moreover, if one ventures further into the realm of speculation, recent computer science results have shown that if one can entangle two parties trying to prove something to another, what they can convincingly prove expands to unimaginable heights, including halting problems. While this celebrated result,  $MIP^* = RE$ ,<sup>82</sup> has not yet been directly connected with the physical world, it is tantalizing to imagine how it might reflect on the power of physical experiments empowered by entanglement to reveal the mysteries of the universe. The question, of course, becomes where to find such useful quantum data and how to effectively get it into the quantum computer.

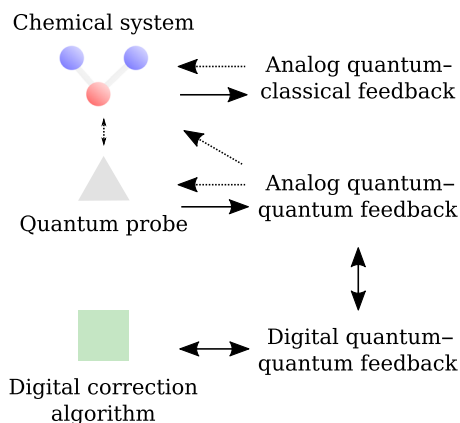
Chemistry has long benefited from the power of quantum sensors in nuclear magnetic resonance experiments. While the most basic forms of these experiments are based on ensemble measurements at an effectively high temperature, they offer incredible insights and are fundamentally quantum. Interestingly, the form of these basic experiments can be closely identified with the so-called one clean qubit model or DQC1 and this has inspired some recent work in pursuing advantages with nuclear magnetic resonance (NMR).<sup>83</sup> However, as NMR technology advances with techniques such as hyper-polarization,<sup>84</sup> multi-dimensional methods,<sup>85</sup>

spatial resolution,<sup>86</sup> and zero-field techniques,<sup>87</sup> we appear to be accessing more and more pure quantum states with increasing control.<sup>88</sup> In addition, the design of more advanced molecular quantum sensors is an active research area.<sup>89</sup> With sufficient advances in transduction techniques and developments in quantum error correction, it may become possible to load data from molecules directly into quantum computers and perform manipulations that are provably challenging for a classical device even with unbounded compute time to replicate due to the query separations we mention above. This could enable faster chemical identification, unprecedented accuracy in examining quantum effects in molecules, or fundamentally new techniques for the control and manipulation of chemical states.

## V. UNTAPPED RESOURCES IN QUANTUM COMPUTER SCIENCE

The discussion thus far has largely centered on the role of quantum computer science or quantum computers themselves in new approaches to understanding chemistry. However, it is also interesting to ask more generally what theoretical tools from quantum information science with potential to impact computational chemistry remain untapped. For example, techniques in tensor networks that received much of their theoretical development in quantum information have now begun to inform the ansatz for correlated wavefunctions. Both matrix product states and tensor network states coupled with an entanglement perspective on strong correlation have led to a wealth of theoretical developments.<sup>90,91</sup>

An area of remarkable theoretical development that remains relatively unknown in the chemical community is that of digital quantum error correction.<sup>16,92</sup> Perhaps for a chemist, the best way to understand these methods is in relation to the existing and well developed body of research of quantum control of chemical reactions and chemistry.<sup>93–95</sup> We depict aspects of this relationship in Fig. 4. In quantum control, a signal, often through a laser excitation or similar, may be used to control and measure the system through feedback; however, these measurements are often destructive with regard to quantum states in the system, and we refer to this as analog quantum-classical feedback control. Through the use of a separate quantum probe or ancilla system, one can make measurements of a system that actually generate or stabilize entanglement in the quantum system or an analog quantum–quantum scheme. Taking this one step further, one may use multiple probes and digital algorithms that adaptively change based on measurement results to pump entropy out much faster than traditional thermalization would allow. It turns out that the exponentially long suppression of transport out of collection of desired quantum states is reliant on an abstract digital perspective of the states. These constructions are formalized in the theory of digital quantum error correction, where it has been shown that with only spatially local measurements and classical computation, it can be possible to stabilize exotic quantum states for essentially arbitrary lengths of time with asymptotically modest resources. While quantum control has a rich history in quantum chemistry and, indeed, many of the developments in reducing errors in quantum experiments<sup>93,96</sup> have relationships to techniques in chemistry, the full digital perspective remains relatively unexplored. If such techniques could be ported more directly into chemistry, stabilizing electronic or excitonic states could lead to

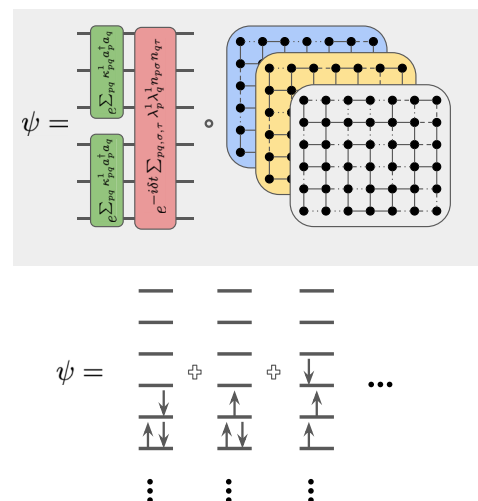


**FIG. 4.** Digital quantum feedback control inspired by quantum error correction. Quantum control of chemical reactions is a well-studied field, and many aspects of it are mirrored in the control of qubit systems. The dotted lines denote quantum interaction, while the solid lines imply classical information exchange. While direct measurement of a system through typical analog signals destroys entanglement in a quantum system (quantum-classical), the use of a quantum probe or ancilla system can allow for measurements that preserve or create entanglement in the target system reliably (quantum-quantum). While this paradigm is already powerful, the tools of discrete computer science allow one to use such interactions in combinations with algorithms that can be imagined as sophisticated Maxwell's demons to pump entropy out of the system at incredible rates far exceeding those of typical thermalization to stabilize exotic quantum states, such as electronic states, for time scales that could, in principle, extend for years, using only local measurements and feedback. These tools are developed at length in the field of quantum error correction, and we imagine here how those tools might eventually integrate with chemical systems. Such constructions may pave the wave for novel excitonic or energy transport designs.

novel reactive methods or wildly improved energy transfer efficiencies. While this remains in the realm of pure speculation perhaps probe systems that can be interacted via future versions of tunneling microscopes with near single electron resolution coupled to catalytically active sites on metal surfaces could offer some of the first places to explore these ideas.

The development of quantum error correction was also closely linked to the development of simulation for a particular class of quantum states, known as stabilizer states. These states can exhibit essentially maximal entanglement across a system yet are incredibly efficient to simulate classically, often scaling like  $N^2$  or  $N^3$  in the number of orbitals.<sup>97,98</sup> While they are discrete in nature, requiring some novel optimization techniques such as reinforcement learning, if coupled with simple orbital rotations, these could form yet another powerful ansatz for efficient classical electronic structure methods informed by quantum information that has yet to be exploited. We depict this connection in Fig. 5.

While the existence of a powerful ansatz that has not yet been explored is perhaps strong enough motivation by itself, the link to quantum error correction provides even more fuel for this interesting avenue of inquiry. Many of the most challenging bio-metallic enzymes, such as nitrogenase, are characterized by an active site that is difficult to describe with traditional methods due to static correlation related to quasi-degeneracy of electronic configurations. Quantum error correction takes advantage of strong entanglement



**FIG. 5.** Contrasting the state vs stabilizer and rotation viewpoint for quasi-degenerate systems. The top shows a representation of an electronic state built from rotations applied to a stabilizer state, while the bottom depicts a traditional configuration based view. Electronic catalysts often exhibit quasi-degeneracy, making them challenging to treat with single reference methods. The stabilizer formalism from quantum error correction, combined with efficient single particle rotations, may offer a compact way to both simulate and analyze such situations, as the use of degeneracy to protect quantum states is a core principle of quantum error correction. Graph state representations of stabilizer states may offer connections to chemical bonding and electron correlation theory, all while being computationally efficient to simulate and analyze classically.

and engineered degeneracy as a mechanism for protecting information. If nature were to follow a similar path, perhaps the electronic near-degeneracy and resulting entanglement in these systems could be an avenue for protecting a coherent reaction pathway similar to the approach of quantum error correction. In addition, many modern quantum error correcting codes take advantage of topological protection, and topological effects have been studied in the context of electronic degeneracies near conical intersection.<sup>99</sup> While natural systems do not have access to an optimal decoder to remove excess entropy, the notion of self-correcting memories where systems exhibit such behavior more naturally draws an enticing connection.<sup>100</sup> In such a case, the stabilizer formalism for representing these ground states would be especially apt, efficiently manipulating untold numbers of determinants through non-obvious symmetries.<sup>101</sup> This could offer an electronic (non-point group) symmetry perspective view on important intermediate or catalytic electronic states. Such statements are presently at the level of conjecture, but it represents another example of a potential way in which quantum computation and chemistry may unite. We depict this Perspective in Fig. 5.

## VI. OUTLOOK

As quantum technology advances, so does our understanding of what any computational device can accomplish. In this Perspective, we have explored the ways in which results from quantum computer science may impact our view and approach toward computational chemistry. On the one hand, quantum computers are believed

to offer an exponential advantage over classical computers in direct simulation of some quantum chemical theories, making certain tasks that previously seemed impossible into ones that should be relatively routine. On the other hand, we saw that even quantum computers have limits, crystallized through the no-fast-forwarding theorem and strong results on the undecidability of physical processes with quantum computers.

In highlighting the ways some of these speed limits are broken, we were led to a framework where learning from natural data is fundamentally different from rote computation. This viewpoint captures the ability of reduced theories to offer meaningful predictions further than the time scales of experiments, and also offers a strategy for dealing with some of the hardest problems in chemical theory.

It has now been decisively shown that classical models empowered by data from quantum computation (including both nature and engineered computers) are more powerful than traditional computation without data, assuming only that quantum computers may perform some tasks faster than classical computers. While data from quantum devices will play an important role in their future interactions with chemistry, the ability to feed quantum data directly into quantum computers offers more power still. The interplay of quantum computers with more advanced quantum sensors may offer untold possibilities. Finally, we expect the interplay between quantum information theory and chemical theory to continue, where import of ideas from domains such as digital quantum error correction may open new avenues of research.

In looking forward, we see a bright future for the ways in which quantum technology may advance the study of chemistry. While it is true that the direct simulation abilities of quantum computers will offer amazing value, we have argued here that this is just the tip of the iceberg. Our ability to address some of the hardest problems in chemistry and nature through data and learning will only be bolstered by stronger quantum technology, and that future is on its way.

## ACKNOWLEDGMENTS

We thank Nathan Wiebe for helpful discussions and feedback on the draft.

## AUTHOR DECLARATIONS

### Conflict of Interest

The authors have no conflicts to disclose.

## DATA AVAILABILITY

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

## REFERENCES

- R. P. Feynman, "Simulating physics with computers," *Int. J. Theor. Phys.* **21**, 467 (1982).
- D. Deutsch, "Quantum theory, the Church–Turing principle and the universal quantum computer," *Proc. R. Soc. London, Ser. A* **400**, 97 (1985).
- P. W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (IEEE, 1994), pp. 124–134.
- A. Ekert and R. Jozsa, "Quantum computation and Shor's factoring algorithm," *Rev. Mod. Phys.* **68**, 733 (1996).
- A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, "Simulated quantum computation of molecular energies," *Science* **309**, 1704 (2005).
- D. Wecker, B. Bauer, B. K. Clark, M. B. Hastings, and M. Troyer, "Gate-count estimates for performing quantum chemistry on small quantum computers," *Phys. Rev. A* **90**, 022305 (2014).
- J. R. McClean, R. Babbush, P. J. Love, and A. Aspuru-Guzik, "Exploiting locality in quantum computation for quantum chemistry," *J. Phys. Chem. Lett.* **5**, 4368 (2014).
- R. Babbush, D. W. Berry, I. D. Kivlichan, A. Y. Wei, P. J. Love, and A. Aspuru-Guzik, "Exponentially more precise quantum simulation of fermions in second quantization," *New J. Phys.* **18**, 033032 (2016).
- R. Babbush, D. W. Berry, J. R. McClean, and H. Neven, "Quantum simulation of chemistry with sublinear scaling in basis size," *npj Quantum Inf.* **5**, 92 (2019).
- G. A. Quantum *et al.*, "Hartree-Fock on a superconducting qubit quantum computer," *Science* **369**, 1084 (2020).
- F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell *et al.*, "Quantum supremacy using a programmable superconducting processor," *Nature* **574**, 505 (2019).
- B. P. Lanyon, J. D. Whitfield, G. G. Gillett, M. E. Goggin, M. P. Almeida, I. Kassal, J. D. Biamonte, M. Mohseni, B. J. Powell, M. Barbieri *et al.*, "Towards quantum chemistry on a quantum computer," *Nat. Chem.* **2**, 106 (2010).
- A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, "A variational eigenvalue solver on a photonic quantum processor," *Nat. Commun.* **5**, 4213 (2014).
- P. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding *et al.*, "Scalable quantum simulation of molecular energies," *Phys. Rev. X* **6**, 031007 (2016).
- A. R. Calderbank and P. W. Shor, "Good quantum error-correcting codes exist," *Phys. Rev. A* **54**, 1098 (1996).
- D. Gottesman, "Stabilizer codes and quantum error correction," [arXiv:quant-ph/9705052](https://arxiv.org/abs/quant-ph/9705052) (1997).
- J. R. McClean, M. E. Kimchi-Schwartz, J. Carter, and W. A. De Jong, "Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states," *Phys. Rev. A* **95**, 042308 (2017).
- W. J. Huggins, S. McArdle, T. E. O'Brien, J. Lee, N. C. Rubin, S. Boixo, K. B. Whaley, R. Babbush, and J. R. McClean, "Virtual distillation for quantum error mitigation," *Phys. Rev. X* (to be published).
- J. Lee, D. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, "Even more efficient quantum computations of chemistry through tensor hypercontraction," *Quantum* **2**, 030305 (2021), <https://journals.aps.org/prxquantum/abstract/10.1103/PRXQuantum.2.030305>.
- N. Schuch and F. Verstraete, "Computational complexity of interacting electrons and fundamental limitations of density functional theory," *Nat. Phys.* **5**, 732 (2009).
- J. D. Whitfield, P. J. Love, and A. Aspuru-Guzik, "Computational complexity in electronic structure," *Phys. Chem. Chem. Phys.* **15**, 397 (2013).
- B. O'Gorman, S. Irani, J. Whitfield, and B. Fefferman, "Electronic structure in a fixed basis is QMA-complete," [arXiv:2103.08215](https://arxiv.org/abs/2103.08215) (2021).
- D. A. Reed, R. Bajcsy, M. A. Fernandez, J.-M. Griffiths, R. D. Mott, J. Dongarra, C. R. Johnson, A. S. Inouye, W. Miner, M. K. Matzke *et al.*, "Computational science: Ensuring America's competitiveness," Technical Report, President's Information Technology Advisory Committee, Arlington, VA, 2005; available at <https://apps.dtic.mil/sti/citations/ADA462840>.
- E. Knill, G. Ortiz, and R. D. Somma, "Optimal quantum measurements of expectation values of observables," *Phys. Rev. A* **75**, 012328 (2007).
- T. Helgaker, P. Jorgensen, and J. Olsen, *Molecular Electronic-Structure Theory* (John Wiley & Sons, 2014).
- B. Sanchez-Lengeling and A. Aspuru-Guzik, "Inverse molecular design using machine learning: Generative models for matter engineering," *Science* **361**, 360 (2018).
- I. Kassal, S. P. Jordan, P. J. Love, M. Mohseni, and A. Aspuru-Guzik, "Polynomial-time quantum algorithm for the simulation of chemical dynamics," *Proc. Natl. Acad. Sci. U. S. A.* **105**, 18681 (2008).

- <sup>28</sup>Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya *et al.*, “Quantum chemistry in the age of quantum computing,” *Chem. Rev.* **119**, 10856 (2019).
- <sup>29</sup>S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, “Quantum computational chemistry,” *Rev. Mod. Phys.* **92**, 015003 (2020).
- <sup>30</sup>V. von Burg, G. H. Low, T. Häner, D. S. Steiger, M. Reiher, M. Roetteler, and M. Troyer, “Quantum computing enhanced computational catalysis,” *Phys. Rev. Research* **3**, 033055 (2021); available at <https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.3.033055>.
- <sup>31</sup>J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, “The theory of variational hybrid quantum-classical algorithms,” *New J. Phys.* **18**, 023023 (2016).
- <sup>32</sup>P. K. Barkoutsos, F. Gkritis, P. J. Ollitrault, I. O. Sokolov, S. Woerner, and I. Tavernelli, “Quantum algorithm for alchemical optimization in material design,” *Chem. Sci.* **12**, 4345 (2021).
- <sup>33</sup>R. Babbush, J. R. McClean, M. Newman, C. Gidney, S. Boixo, and H. Neven, “Focus beyond quadratic speedups for error-corrected quantum advantage,” *Phys. Rev. X Quantum* **2**, 010103 (2021).
- <sup>34</sup>J. K. Nørskov, T. Bligaard, J. Rossmeisl, and C. H. Christensen, “Towards the computational design of solid catalysts,” *Nat. Chem.* **1**, 37 (2009).
- <sup>35</sup>D. Ess, L. Gagliardi, and S. Hammes-Schiffer, “Introduction: Computational design of catalysts from molecules to materials,” *Chem. Rev.* **119**, 6507 (2019).
- <sup>36</sup>G. Knizia, “Intrinsic atomic orbitals: An unbiased bridge between quantum theory and chemical concepts,” *J. Chem. Theory Comput.* **9**, 4834 (2013).
- <sup>37</sup>J. Hachmann, R. Olivares-Amaya, S. Atahan-Evrenk, C. Amador-Bedolla, R. S. Sánchez-Carrera, A. Gold-Parker, L. Vogt, A. M. Brockway, and A. Aspuru-Guzik, “The Harvard clean energy project: Large-scale computational screening and design of organic photovoltaics on the world community grid,” *J. Phys. Chem. Lett.* **2**, 2241 (2011).
- <sup>38</sup>A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder *et al.*, “Commentary: The materials project: A materials genome approach to accelerating materials innovation,” *APL Mater.* **1**, 011002 (2013).
- <sup>39</sup>A. K. Singh, K. Mathew, H. L. Zhuang, and R. G. Hennig, “Computational screening of 2D materials for photocatalysis,” *J. Phys. Chem. Lett.* **6**, 1087 (2015).
- <sup>40</sup>S. L. Kinnings, N. Liu, P. J. Tonge, R. M. Jackson, L. Xie, and P. E. Bourne, “A machine learning-based method to improve docking scoring functions and its application to drug repurposing,” *J. Chem. Inf. Model.* **51**, 408 (2011).
- <sup>41</sup>L.-P. Wang, A. Titov, R. McGibbon, F. Liu, V. S. Pande, and T. J. Martínez, “Discovering chemistry with an *ab initio* nanoreactor,” *Nat. Chem.* **6**, 1044 (2014).
- <sup>42</sup>M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, “Elucidating reaction mechanisms on quantum computers,” *Proc. Natl. Acad. Sci. U. S. A.* **114**, 7555 (2017).
- <sup>43</sup>M. Motta, E. Ye, J. R. McClean, Z. Li, A. J. Minnich, R. Babbush, and G. K. Chan, “Low rank representations for quantum simulation of electronic structure,” *NPJ Quantum Information* **7**, 83 (2021).
- <sup>44</sup>D. W. Berry, C. Gidney, M. Motta, J. R. McClean, and R. Babbush, “Qubitization of arbitrary basis quantum chemistry leveraging sparsity and low rank factorization,” *Quantum* **3**, 208 (2019).
- <sup>45</sup>A. D. Bochevarov, E. F. Valeev, and C. David Sherrill, “The electron and nuclear orbitals model: Current challenges and future prospects,” *Mol. Phys.* **102**, 111 (2004).
- <sup>46</sup>T. J. Martínez, “Ab initio reactive computer aided molecular design,” *Acc. Chem. Res.* **50**, 652 (2017).
- <sup>47</sup>D. R. Weinberg, C. J. Gagliardi, J. F. Hull, C. F. Murphy, C. A. Kent, B. C. Westlake, A. Paul, D. H. Ess, D. G. McCafferty, and T. J. Meyer, “Proton-coupled electron transfer,” *Chem. Rev.* **112**, 4016 (2012).
- <sup>48</sup>S. Gu, R. D. Somma, and B. Şahinoğlu, “Fast-forwarding quantum evolution,” *arXiv:2105.07304* (2021).
- <sup>49</sup>R. Greenlaw, H. J. Hoover, W. L. Ruzzo *et al.*, *Limits to Parallel Computation: P-Completeness Theory* (Oxford University Press on Demand, 1995).
- <sup>50</sup>M. Coudron and S. Menda, “Computations with greater quantum depth are strictly more powerful (relative to an oracle),” in *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing* (ACM, 2020), pp. 889–901.
- <sup>51</sup>D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, “Efficient quantum algorithms for simulating sparse Hamiltonians,” *Commun. Math. Phys.* **270**, 359 (2007).
- <sup>52</sup>Y. Atia and D. Aharonov, “Fast-forwarding of Hamiltonians and exponentially precise measurements,” *Nat. Commun.* **8**, 1572 (2017).
- <sup>53</sup>C. Dellago, P. G. Bolhuis, F. S. Csajka, and D. Chandler, “Transition path sampling and the calculation of rate constants,” *J. Chem. Phys.* **108**, 1964 (1998).
- <sup>54</sup>J. C. Phillips, R. Braun, W. Wang, J. Gumbart, E. Tajkhorshid, E. Villa, C. Chipot, R. D. Skeel, L. Kalé, and K. Schulten, “Scalable molecular dynamics with NAMD,” *J. Comput. Chem.* **26**, 1781 (2005).
- <sup>55</sup>J. Bucklew, *Introduction to Rare Event Simulation* (Springer Science & Business Media, 2013).
- <sup>56</sup>B. B. Machta, R. Chachra, M. K. Transtrum, and J. P. Sethna, “Parameter space compression underlies emergent theories and predictive models,” *Science* **342**, 604 (2013).
- <sup>57</sup>M. K. Transtrum, B. B. Machta, K. S. Brown, B. C. Daniels, C. R. Myers, and J. P. Sethna, “Perspective: Slowness and emergent theories in physics, biology, and beyond,” *J. Chem. Phys.* **143**, 010901 (2015).
- <sup>58</sup>N. Shiraishi and K. Matsumoto, “Undecidability in quantum thermalization,” *Nat. Comm.* **12**, 5084 (2021).
- <sup>59</sup>J. v. Neumann, “Proof of the quasi-ergodic hypothesis,” *Proc. Natl. Acad. Sci. U. S. A.* **18**, 70 (1932).
- <sup>60</sup>E. H. Kennard *et al.*, *Kinetic Theory of Gases* (McGraw-Hill, New York, 1938), Vol. 287.
- <sup>61</sup>L. Boltzmann, *Lectures on Gas Theory* (Courier Corporation, 2012).
- <sup>62</sup>T. S. Cubitt, D. Perez-Garcia, and M. M. Wolf, “Undecidability of the spectral gap,” *Nature* **528**, 207 (2015).
- <sup>63</sup>P. W. Humphreys, “More is different. . . sometimes: Ising models, emergence, and undecidability,” in *Why More Is Different* (Springer, 2015), pp. 137–152.
- <sup>64</sup>J. Bausch, T. S. Cubitt, A. Lucia, and D. Perez-Garcia, “Undecidability of the spectral gap in one dimension,” *Phys. Rev. X* **10**, 031038 (2020).
- <sup>65</sup>A. Turing, *Halting Problem* (London Mathematical Society, 1936).
- <sup>66</sup>R. G. Palmer, “Broken ergodicity,” *Adv. Phys.* **31**, 669 (1982).
- <sup>67</sup>S. Chikazumi and C. D. Graham, *Physics of Ferromagnetism* (Oxford University Press on Demand, 2009), Vol. 2e, p. 94.
- <sup>68</sup>I. Prigogine and G. Nicolis, “On symmetry-breaking instabilities in dissipative systems,” *J. Chem. Phys.* **46**, 3542 (1967).
- <sup>69</sup>H.-O. Georgii, *Gibbs Measures and Phase Transitions* (Walter de Gruyter, 2011), Vol. 9.
- <sup>70</sup>C. Moore, “Unpredictability and undecidability in dynamical systems,” *Phys. Rev. Lett.* **64**, 2354 (1990).
- <sup>71</sup>G. J. Chaitin, “A theory of program size formally identical to information theory,” *J. ACM* **22**, 329 (1975).
- <sup>72</sup>R. Freund, L. Kari, and G. Păun, “DNA computing based on splicing: The existence of universal computers,” *Theory Comput. Syst.* **32**, 69 (1999).
- <sup>73</sup>G. Paun, G. Rozenberg, and A. Salomaa, *DNA Computing: New Computing Paradigms* (Springer Science & Business Media, 2005).
- <sup>74</sup>S. Arora and B. Barak, *Computational Complexity: A Modern Approach* (Cambridge University Press, 2009).
- <sup>75</sup>H.-Y. Huang, M. Broughton, M. Mohseni, R. Babbush, S. Boixo, H. Neven, and J. R. McClean, “Power of data in quantum machine learning,” *Nature Communications* **12**, 2631 (2021).
- <sup>76</sup>O. A. von Lilienfeld and K. Burke, “Retrospective on a decade of machine learning for chemical discovery,” *Nat. Commun.* **11**, 4895 (2020).
- <sup>77</sup>J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, “Neural message passing for quantum chemistry,” in *International Conference on Machine Learning* (PMLR, 2017), pp. 1263–1272.
- <sup>78</sup>M. Aykol, V. I. Hegde, L. Hung, S. Suram, P. Herring, C. Wolverton, and J. S. Hummelshøj, “Network analysis of synthesizable materials discovery,” *Nat. Commun.* **10**, 2018 (2019).
- <sup>79</sup>S. Aaronson, “Guest column: NP-complete problems and physical reality,” *ACM SIGACT News* **36**, 30 (2005).
- <sup>80</sup>H.-Y. Huang, R. Kueng, and J. Preskill, “Information-theoretic bounds on quantum advantage in machine learning,” *Phys. Rev. Lett.* **126**, 190505 (2021).



- <sup>81</sup>D. Aharonov, J. Cotler, and X.-L. Qi, "Quantum algorithmic measurement," [arXiv:2101.04634 \[quant-ph\]](#) (2021).
- <sup>82</sup>Z. Ji, A. Natarajan, T. Vidick, J. Wright, and H. Yuen, "MIP\* = RE," [arXiv:2001.04383](#) (2020).
- <sup>83</sup>D. Sels, H. Dashti, S. Mora, O. Demler, and E. Demler, "Quantum approximate Bayesian computation for NMR model inference," *Nat. Mach. Intell.* **2**, 396 (2020).
- <sup>84</sup>J. C. Leawoods, D. A. Yablonskiy, B. Saam, D. S. Gierada, and M. S. Conradi, "Hyperpolarized <sup>3</sup>He gas production and MR imaging of the lung," *Concepts Magn. Reson.* **13**, 277 (2001).
- <sup>85</sup>R. R. Ernst and W. A. Anderson, "Application of Fourier transform spectroscopy to magnetic resonance," *Rev. Sci. Instrum.* **37**, 93 (1966).
- <sup>86</sup>S. J. DeVience, L. M. Pham, I. Lovchinsky, A. O. Sushkov, N. Bar-Gill, C. Belthangady, F. Casola, M. Corbett, H. Zhang, M. Lukin *et al.*, "Nanoscale NMR spectroscopy and imaging of multiple nuclear species," *Nat. Nanotechnol.* **10**, 129 (2015).
- <sup>87</sup>D. B. Zax, A. Bielecki, K. W. Zilm, A. Pines, and D. P. Weitekamp, "Zero field NMR and NQR," *J. Chem. Phys.* **83**, 4877 (1985).
- <sup>88</sup>Y. Xia, W. Li, W. Clark, D. Hart, Q. Zhuang, and Z. Zhang, "Demonstration of a reconfigurable entangled radio-frequency photonic sensor network," *Phys. Rev. Lett.* **124**, 150502 (2020).
- <sup>89</sup>C.-J. Yu, S. von Kugelgen, D. W. Laorenza, and D. E. Freedman, "A molecular approach to quantum sensing," *ACS Cent. Sci.* **7**, 712 (2021).
- <sup>90</sup>S. Sharma, K. Sivalingam, F. Neese, and G. K.-L. Chan, "Low-energy spectrum of iron-sulfur clusters directly from many-particle quantum mechanics," *Nat. Chem.* **6**, 927 (2014).
- <sup>91</sup>G. Barcza, Ö. Legeza, K. H. Marti, and M. Reiher, "Quantum-information analysis of electronic states of different molecular structures," *Phys. Rev. A* **83**, 012508 (2011).
- <sup>92</sup>S. A. Ocko, X. Chen, B. Zeng, B. Yoshida, Z. Ji, M. B. Ruskai, and I. L. Chuang, "Quantum codes give counterexamples to the unique preimage conjecture of the *N*-representability problem," *Phys. Rev. Lett.* **106**, 110501 (2011).
- <sup>93</sup>H. Rabitz, R. de Vivie-Riedle, M. Motzkus, and K. Kompa, "Whither the future of controlling quantum phenomena?," *Science* **288**, 824 (2000).
- <sup>94</sup>C. Brif, R. Chakrabarti, and H. Rabitz, "Control of quantum phenomena: Past, present and future," *New J. Phys.* **12**, 075008 (2010).
- <sup>95</sup>M. Shapiro and P. Brumer, "Quantum control of chemical reactions," *J. Chem. Soc., Faraday Trans.* **93**, 1263 (1997).
- <sup>96</sup>L. Viola, E. Knill, and S. Lloyd, "Dynamical decoupling of open quantum systems," *Phys. Rev. Lett.* **82**, 2417 (1999).
- <sup>97</sup>D. Gottesman, "The Heisenberg representation of quantum computers," in *Group22: Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics*, edited by S. P. Corney, R. Delbourgo, and P. D. Jarvis (Cambridge, MA, International Press, 1999), pp. 32–43.
- <sup>98</sup>S. Aaronson and D. Gottesman, "Improved simulation of stabilizer circuits," *Phys. Rev. A* **70**, 052328 (2004).
- <sup>99</sup>D. R. Yarkony, "Diabolical conical intersections," *Rev. Mod. Phys.* **68**, 985 (1996).
- <sup>100</sup>B. J. Brown, D. Loss, J. K. Pachos, C. N. Self, and J. R. Wootton, "Quantum memories at finite temperature," *Rev. Mod. Phys.* **88**, 045005 (2016).
- <sup>101</sup>H. J. García, I. L. Markov, and A. W. Cross, "On the geometry of stabilizer states," *Quantum Inf. Comput.* **14**, 683–720 (2014).