This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1147/JRD.2018.2888987, IBM Journal of Research and Development

Potential of quantum computing for drug discovery

Yudong Cao[1,2,*], Jhonathan Romero[1,2] and Alán Aspuru-Guzik[1,2,3,4,5]

***Abstract*— Quantum computing has rapidly advanced in recent years due to substantial development in both hardware and algorithms. These advances are carrying quantum computers closer to their impending commercial utility. Drug discovery is a promising area of application which will find a number of uses for these new machines. As a prominent example, quantum simulation will enable faster and more accurate characterizations of molecular systems than existing quantum chemistry methods. Furthermore, algorithmic developments in quantum machine learning offer interesting alternatives to classical machine learning techniques, which may also be useful for the biochemical efforts involved in early phases of drug discovery. Meanwhile, quantum hardware is scaling up rapidly into a regime where an exact simulation is difficult even using the world's largest supercomputers. We review how these recent advances can shift the paradigm with which one thinks about drug discovery, focusing on both the promises and caveats associated with each development. In particular, we highlight how hybrid quantum-classical approaches to quantum simulation and quantum machine learning could yield substantial progress using noisy-intermediate scale quantum devices, while fault-tolerant, error corrected quantum computers are still in their development phase.**

Drug discovery is the process of developing a drug from an initial hypothesis to a fully commercialized product. This process can often take more than a decade and billions of dollars in expenditure before a molecule can be recognized as a drug [1]. A significant portion of these resources is invested in the identification of molecules that exhibit significant medicinal activity against a disease, usually referred to as a *hit*. Most of the research in drug discovery focuses on hits of low molecular weight (<900 Daltons, with sizes of 1nm or less [2]), which constitute around 78% of the drug market [3]. In this case, the medicinal activity of a particular drug candidate or *ligand* is associated to its ability to bind to a biological *target*, usually a protein, whose activity regulates the metabolism of the disease. Typically, the first stage in the discovery process is to generate a library of potential drug candidates that is subsequently screened based on medicinal activity to identify hits [4]. Along with the activity, other factors that determine the efficacy and potency of the hits, such as the absorption, distribution, metabolism, excretion and toxicity (ADMET) profile, among other pharmacokinetic properties, are optimized to produce a smaller set of better candidates called *lead compounds* [5]. Further screening and optimization generally delivers a small set of leads which proceed through the stages of drug development and clinical trials before one of them becomes a viable commercial product.

Traditionally, the search of hits was accomplished by high-throughput screening (HTS) on large molecular libraries using *in vitro* activity experiments. These searches generally have low hit rates and required the synthesis of a large number of compounds, which in turn demanded a significant investment of resources and time [5]. This approach was completely transformed by the advent of commercial computers in the 70s and 80s, which enabled computational chemistry and statistical analysis, among other tools [6], to accelerate HTS, improve the hit rate and increase the quality of the leads obtained in the process [5]. The increase in computational power and the improvement of computational chemistry techniques fostered the practice of *Computer-Aided Drug Design* (CADD), which constitutes a significant portion of the drug discovery pipeline today. The ultimate goal of CADD is to answer the inverse-design question: *what are the best chemical structures associated with a desired therapeutic effect?* [7], [8]. To answer this question accurately, CADD faces two main challenges: (1) the accurate simulation of the interaction of drug candidates with biological targets, and (2) accurate statistical modeling of activities and ADMET profiles based on the available simulated and experimental data. The former is largely constrained by the computational cost of simulating the physics of molecular systems for both small molecules and biological targets. The latter is constrained by the effectiveness of existing statistical techniques.

Quantum computing could potentially shift the paradigm with which one thinks about quantum chemical simulation. By efficiently preparing highly entangled states that are otherwise intractable to describe on classical computers, quantum computers can perform certain important quantum chemistry and machine learning tasks in ways that are beyond the ability of classical computers. Furthermore, efficient manipulation of quantum states also allows for certain linear algebraic operations to be performed far more efficiently than what is possible with classical devices. With these unique abilities, quantum computing promises to deliver efficient and highly accurate solutions to otherwise intractable problems, for instance finding the ground state energy of a molecular system [9]. As we will discuss later in detail later in Section II, a common method for treating electronic structure calculation on a quantum computer is via second quantization, where an electronic state over N spin orbitals is

[1]Department of Chemistry and Chemical Biology, Harvard University, Cambridge MA 02138, USA
[*]Email: yudong@zapatacomputing.com
[2]Zapata Computing Inc., Cambridge, MA 02138, USA
[3]Canadian Institute for Advanced Research, Toronto, Ontario M5G 1Z8, Canada
[4]Department of Computer Science, University of Toronto, Toronto, Ontario M5G 1Z8, Canada
[5]Department of Chemistry, University of Toronto, Toronto, Ontario M5G 1Z8, Canada

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1147/JRD.2018.2888987, IBM Journal of Research and Development

represented using N qubits - one qubit for each spin orbital. In the coming years we are anticipating quantum devices with $N > 50$ qubits [10], [11], making it possible to map onto a quantum computer problems whose exact solution (say via exact Hamiltonian diagonalization) is beyond current classical computation.

Quantum machine learning is also a rapidly emerging field exploring how quantum computers can perform machine learning tasks with improved performance over classical computers [12]. As we will discuss in Section II-B, there are plausible reasons to believe that quantum computers may enable solutions to machine learning tasks that are beyond classical computation. Pinpointing the precise regimes of quantum advantage is a main mission of the field of quantum machine learning.

In this article we review developments in quantum computing relevant to drug discovery through quantum chemistry and machine learning, outlining the promises as well as caveats. The paper is organized as follows: first we describe the general pipeline of CADD and some of the methodologies employed in the industry and their challenges. Second, we outline some of the latest quantum computing algorithms that we consider relevant for CADD, namely quantum simulation and quantum machine learning. Finally, we share our perspective on how these methods could benefit CADD by addressing some of its biggest challenges.

The purpose of this perspective is to initiate a mutually beneficial cross-disciplinary discussion and collaboration between the fields of CADD and quantum computing. For the quantum computing community such dialogue will help to outline the practically useful regimes where quantum computers may have an advantage over classical counterparts. For the drug discovery community our hope is to bring an alternative perspective on classical computing for solving some of the crucial computational problems which arise in practice. Our approach is by no means exhaustive, and for more in-depth discussion of specific technical subjects, the reader is encouraged to refer to the relevant citations.

I. Overview of computational methods in drug discovery

At the risk of gross simplification, we summarize the overall drug discovery process as shown in Figure 1a. The usual drug discovery pipeline requires the identification and characterization of a suitable biological target, which can be effectively proved to intervene in the mechanism of disease. This step often requires intense experimentation as well as extensive statistical analysis of the collected data. Once a biological target is in place, the next step is the search for hits, which usually involves extensive biological and virtual screening over libraries of molecules, or the generation of completely new compounds (*de novo* design), that must be synthesized and tested. The group of hits collected on this stage undergoes further optimization of the pharmacokinetics and ADMET properties, involving a combination of biological and *in silico* tests, to generate the final group of leads. These stages, going from target identification to lead optimization, benefit the most from CADD techniques [13], [14]. The subsequent steps in the drug discovery pipeline, which involve clinical studies in animals and humans prior to the Federal Drug Administration (FDA) review and approval are less intensive in the use of CADD tools, but might require further rounds of lead optimization. The final success of a drug discovery campaign depends, to a great extent, on the quality of the CADD approaches applied in the early stages.

CADD approaches employed on the stages of hit search, lead discovery and lead optimization are generally classified into two main categories [5]: *structured-based* and *ligand-based* (Figure 1b). Structured-based CADD relies on knowledge of the target protein 3D structure to predict the ability of a candidate to bind to the target, whereas ligand-based CADD employs information of known active and inactive molecules to predict the activity of new candidates. Structure-based CADD is preferred over ligand-based if the structural information of the biological target is available. This information is usually obtained experimentally using NMR spectroscopy and X-ray crystallography studies on crystallized protein [5]. Predicting the protein structure from the knowledge of the amino-acid sequence requires simulating the protein folding process, which is so far out of reach except for small peptides and fast folders [15]; however, in the absence of experimental structures, it is still possible to approximate the 3D structure of an unknown target protein by comparing its sequence with related known proteins, a process known as comparative modeling [16]. Along with the structure, it is necessary to characterize the target by identifying the binding (active) sites, that are responsible for the biological activity and where the potential drug candidate (*ligand*) is expected to bind.

Assuming that a model for the target structure is available, structured-based CADD approaches attempt to find suitable drug candidates by analyzing the interaction between the candidate (ligand) and the biological target, generally a protein; therefore, most CADD approaches require: (1) determining the *pose* or conformation of the ligand that fits best the binding site of the target, and (2) assigning a numerical score that express the strength of the interaction of the ligand-target complex [17]. The process of finding the best conformation is generally called *docking* and the process of computing the affinity is referred as *scoring* [5]. These procedures are generally intertwined since docking requires a score function that ranks different conformations according to their ability to form bound ligand-protein complexes. Extensive sampling of conformations is often required in structured-based CADD approaches to account for the mobility of protein and ligands in biological conditions (aqueous solutions at room temperature) [18].

When no information of the 3D tertiary structure of the protein can be obtained, ligand-based CADD is the main tool. In this case, the selection of the candidates proceeds by comparison of the structures with a set of known active ligands using molecular similarity indexes and by evaluation of the activity using a *quantitative structure-activity relation* (QSAR) model [5]. QSARs are mathematical models that ex-

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1147/JRD.2018.2888987, IBM Journal of Research and Development

press the activity or in general any other property of interest, as a function of a set of molecular descriptors [19]. Typical descriptors employed in ligand-based CADD encode a variety of chemical information, including molecular weight, geometry, volume, surface areas, ring content, 3D geometrical information, atom types, electronegativities, polarizabilities, molecular symmetry, atom distribution, topological charge indices, functional group composition, aromaticity indices, solvation properties, among others [5]. Both QSAR models and structure comparison require experimental information for the set of molecules employed as either reference for the comparison or training set for QSAR. Consequently, the models employed in ligand-based CADD are often restricted to libraries of candidates that share sufficient similarities with the set of molecules employed as reference [20].

For the rest of this section, we will describe in more detail the different stages of CADD along with some of the structured-based and ligand-based methodologies employed in each of them, as well as some of their limitations and challenges. We place a special focus on *ab initio* methods since those are methods that are most amenable for adapting quantum computing techniques.

A. Target identification and characterization

The initial stage of drug discovery concerns collecting evidence of therapeutic effects in activation or inhibition of certain biological pathway associated to a disease. The biological entity responsible for such response is called the *target*, which is generally a protein. In a broader sense, the term could also refer to the genes or RNA associated to the protein. Ideal targets should be "druggable", meaning that the drug candidate should be able to access the target and effect a biological response that is measurable *in vitro* and *in vivo*. The identification of suitable targets and their corresponding validation by studies of the mechanism of action increases the chance of success during the discovery process and allows one to foresee side effects associated to the modulation of the target [13].

Traditional approaches to target identification employ chemical proteomic techniques such as affinity chromatography, biochemical fractionation, and radioactive ligand binding assays [22]. These methods employ a small molecule with proven activity to isolate the target from a mixture of other proteins. In the case of affinity chromatography, the most widely used approach, the active compound is immobilized in a porous matrix. Subsequently, a solution containing the protein mixture is passed through the matrix, and those proteins that bind to the immobilized active compound are retained. In the final stage, the retained proteins, which correspond to the potential targets, are eluted from the matrix [23]–[25]. The identification of the protein is usually performed via isotopic-labeling and mass spectrometry techniques.

More recently, the improvement of sequencing techniques have enabled the use of genetic screening techniques, where targets are identified by studying the effect of different concentrations of an active compound in a population of cells with different mutations [25], [26]. By identifying which populations are more susceptible to exposition to an active compounds and looking at their mutations, it is possible to infer which proteins are associated to the activity. Another technique called *gene profiling* combines gene expression analysis (e.g. message RNA profiles) with chemical studies to identify targets. This approach is based on the assumption that deleting the genes that codify the target protein should produce the same inhibitory effect of the active compounds. Consequently, the target can be identified by comparing the expression profiles (information of which proteins are synthesized or *expressed*) of the population of mutants with the profiles of populations exposed to the active compound [24]. A similar idea can be applied to identify targets by examining message RNA/protein levels to determine whether they correlate with the manifestation of the disease [13].

The development of systems biology has motivated a more integrative approach to target identification that seeks to identify target associations and their mechanisms within complex biological systems using computational models. Many of these approaches compare available biological data on regulatory networks, molecular pathways and cell phenotypes with the same type of profiles for bioactive molecules and targets in order to find common patterns of biological response and drug activity [27]. The biological data is usually available in a variety of sources that include medical databases, patents, papers and other information available in the web [13], which has motivated the incorporation of data mining strategies in target identification [28], and the use of multiple computational approaches to identify patterns in the data leading to target recognition [29]. Some of these approaches include machine learning methods based on Bayesian inference [30].

Following target identification and validation, the tertiary structure of the target protein must be characterized in order to perform structured-based drug discovery. This can be done using NMR and X-ray crystallography techniques on a protein crystal, which are generally difficult to obtain due to the experimental difficulties of purifying and crystallizing proteins [31]. In the case of targets without known experimental tertiary structures, computational approaches such as threading, comparative modeling and *ab initio* methods can be used to predict 3D structures from the knowledge of the protein aminoacid sequence. Comparative modeling is by far the most common approach for 3D structure prediction. This method generally consists of building the 3D structures of a protein by comparing its sequence with those of proteins whose 3D structures have been experimentally characterized (i.e. 40% similarity). A typical prediction with comparative modeling implies the identification of related proteins that can serve as templates from a database of known structures, usually the Protein Data Bank (PDB) [32]. The sequences of unknown proteins and the templates are aligned and compared, and the geometries of the regions with good alignment are copied. Missing regions require further refinement using relaxation methods based on molecular dynamics, Monte Carlo minimization, or genetic algorithms

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1147/JRD.2018.2888987, IBM Journal of Research and Development

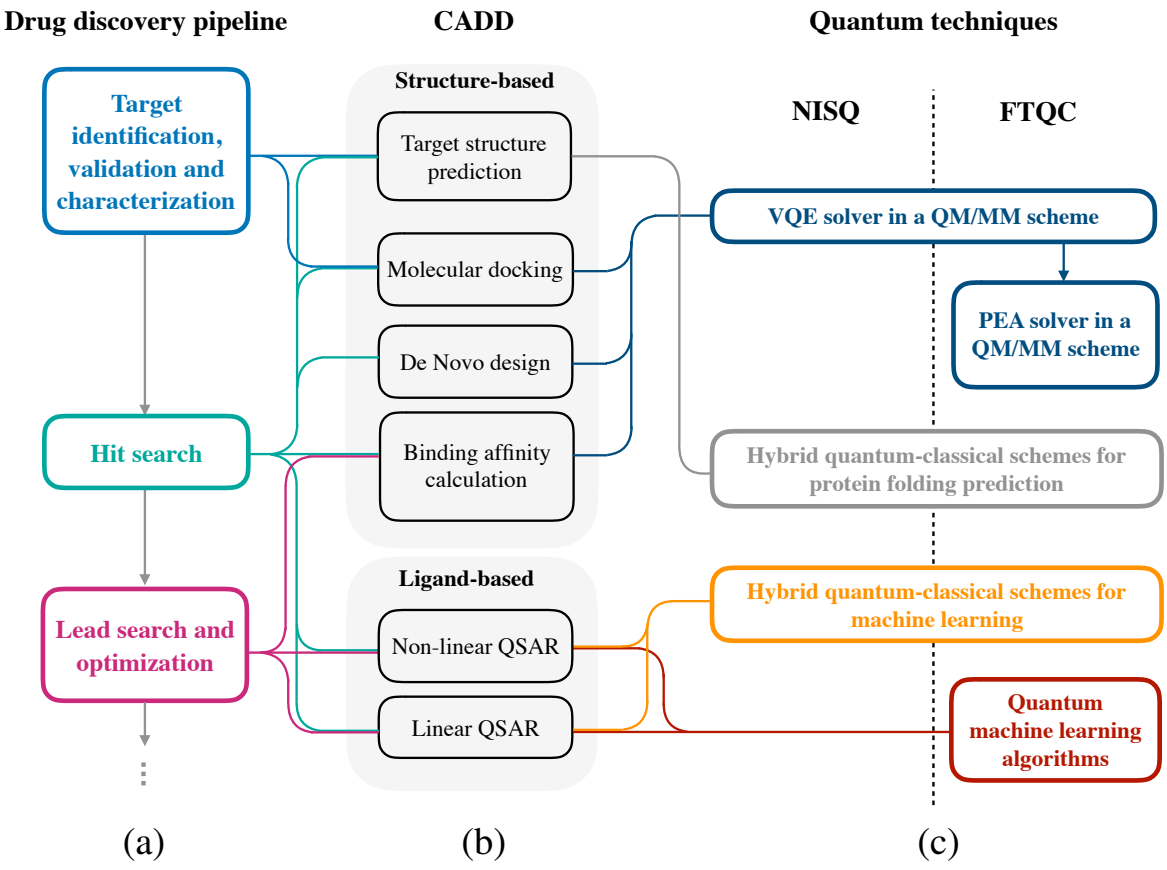

Fig. 1. (a) General workflow of drug discovery process. Here we focus on the early phase where computationally intensive quantum chemical analyses are involved. (b) Components of each stage of drug discovery that heavily involve quantum chemistry or machine learning techniques. (c) Quantum techniques that can be applied to the components listed in (b) and potentially yield an advantage over known classical methods. Here we make the separation between techniques for noisy intermediate scale quantum (NISQ) devices [21] and fault-tolerant quantum computing devices.

[5]. Threading works with a similar principle, but employs comparison with shorter sequences obtained from proteins with less similarity..

In cases where suitable templates are not available for the target of interest, *free modeling* approaches are employed. This term groups knowledge-based approaches for structure prediction and physically motivated methods [16]. Knowledge-based approaches usually assemble a protein structure by using the geometries of small protein fragments extracted from known 3D structures. Purely first principles prediction has been mainly limited to the use of MD simulations with appropriate solvent models to optimize the structure of the protein [16], an approach that is much more computationally demanding than comparative modeling or knowledge-based free methods. The advantage of physical based methods is that they reveal the pathway of protein folding for the unknown structure, but their main bottlenecks are their need for extensive conformational sampling and accurate force-field potentials [33]. In this area, quantum simulation and quantum optimization approaches could become powerful tools by addressing problems such as protein folding, as described in Section III.

A crucial aspect of target characterization is the recognition of binding sites, which are those regions of the protein where the drug candidate is more likely to bind and effect the desirable therapeutic effect [17]. Often, binding sites for small molecules are known for co-crystallized structures of the target or related proteins [5]; however, if the binding sites are unknown or if the purpose is to identify new binding sites for new applications, a few computational programs such as Ligsite, Qsite finder and CASTp, among others [34], [35] can be applied. Most of these methods employ geometrical or atomistic probes to identify places where a small molecule is more likely to bind based on geometrical constraints or on empirical energy functions. Binding site identification could benefit from structured-based techniques such as molecular docking and affinity calculations [17], although this might require extensive conformational search over the surface of the protein, which is computationally more demanding.

B. Hit search

The process of hit search generally involves high-throughput screening (HTS) of a database of candidate compounds. Traditionally, this process has required the synthesis and experimental determination of the activity of the compounds, which is extremely expensive and slow.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1147/JRD.2018.2888987, IBM Journal of Research and Development

Nowadays, the process is accelerated using virtual HTS (vHTS). Different score functions are employed to rank the activity of the candidates depending on whether a structured-based or ligand-based approach is used. Some ligand-based approaches score the candidates based on their similarity with a set of known active compounds. Another option is QSAR, which constructs a statistical model based on experimental information of the activities and chemical information of the ligands. In both approaches, the chemical information is expressed with molecular descriptors that encode physicochemical and structural information of the molecules in a digital format, suitable for comparison. Molecular descriptors can be generated by knowledge-based, graph-theoretical, molecular mechanical or quantum-mechanical methods [36], [37]. Arguably, the most popular descriptors are molecular fingerprints, which encode various molecular properties as predefined bit settings [38], [39]. Other descriptors are computed solely from the 2D or 3D topology of the molecule based on graph-theoretical methods [40].

Comparison-based approaches usually employ molecular fingerprints to compute similarity indexes between molecules, which makes them an efficient tool for vHTS; however, the method suffers from the influence of unnecessary features and is usually limited to a small set of molecules. An alternative to comparison-based methods are QSAR models, which are often applied when a set of compounds with varying degrees of activities is well characterized experimentally and this data can be used to train a model [41]. Consequently, the success of the QSAR will depend on the quality of the initial data along with the selections of descriptors and statistical model. A major drawback is that the applicability of the model is limited to the sampling space of the initial set of compounds, which limits the chemical diversity of the drug candidates obtained in the process. Furthermore, these new candidates are likely to show only small changes in activity with respect to the original compound. For the same reason, QSAR can be a useful tool during the lead optimization phase and is also applied to the prediction of pharmacokinetics or ADMET profiles [42]. While QSAR modeling has many success stories in CADD, its application is generally limited to families of molecules with similar scaffolds [5].

A fundamental aspect of QSAR is the selection of the statistical model. Most QSAR applications involve linear models, such as multivariable linear regression (MLR), principle component analysis (PCA) and partial least square analysis (PLS) [36]. The main advantage of these methods is their computational efficiency, however, they are insufficient to describe the often non-linear relation between biological activity and molecular properties [5]. This fact, along with the increasing amount of activity data available nowadays, has motivated the use of non-linear statistical models for QSAR, most of them employing Artificial Neural Networks (ANN) and other machine learning techniques such as Support Vector Machines (SVM), Random Forest (RF) and Decision Trees (DT) [43]. These approaches have been successfully employed in the prediction of activities [44] and ADMET profiles [45]. Although applications of the deep neural networks to drug design are still exploratory, there are good indications that they could provide new avenues for estimating binding affinity prediction, generating *de novo* drug candidates and predicting ADMET profiles [46]–[49].

Alternatively, structure-based hit search focuses on predicting the protein-ligand geometries and estimating the protein-ligand score functions, which are related with activity. Molecular docking is usually performed for several binding sites on a database of targets. Different docking approaches employ either atomic, surface or grid-based representation of the ligand and target structures [50]. Surface methods describe structures as networks of surfaces with different shapes that mimics the Van Der Waals mappings of the atoms [51]. In this case, docking is performed by aligning the corresponding structures. Grid-based methods simplify the 3D information of the molecule by creating grids to quantify the overlap between the target and the ligand structure. Finally, atomistic models generally rely on force-field simplification of the atomic interactions. The parameters of these force-fields are usually optimized to match quantum mechanical calculations on small molecules, therefore, they are more accurate than grid or surface models. Most molecular docking methods today employ flexible atomistic models, which consider ligand and protein flexibility during the docking process [5]. There are different strategies for sampling the conformations of the ligands and target during docking including systematic enumeration of the confirmations, molecular dynamic simulations, Monte Carlo search algorithms with Metropolis criterion (MCM) and genetic algorithms to find the optimal ligand-protein complexes [5].

Current approaches to generate candidate structures in ligand-based CADD and to sample conformations in structured-based CADD are very efficient and do not constitute a limitation for vHTS. In contrast, the accuracy and efficiency of vHTS is mainly limited by the quality and efficiency of the scoring functions employed in the process [20]. Ideally, scoring functions should be directly related to the binding affinity of the candidate with the target, which is defined as the free energy of the formation of the ligand-target complex from the separate ligand and target molecules. Most scoring functions are derived from statistical models of experimentally characterized ligand-protein complexes or by a physical description of the interactions. The former are further classified into *empirical functions*, where the model employs molecular descriptors and *knowledge-based* functions, where the model employs structural information such as atom-atom distances [5], [52]. Other types of scoring functions are based on molecular-mechanics (MM) calculations (*MM-based*) that employ force fields parameterized using quantum mechanical calculations or directly from experimental information [53]. Finally, *consensus-scoring* functions, which are combinations of the previous types, have recently gained popularity, exhibiting better results in some cases [54]. Despite their computational efficiency and extensive usage, existing scoring functions are far from being infallible and often render false positives during the screening

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1147/JRD.2018.2888987, IBM Journal of Research and Development

process [20].

The main advantage of MM-based scoring functions is their wider range of applicability, compared to empirical and knowledge-based scoring functions that can be only used for reduced molecular spaces for which experimental data is available; however, the accuracy of MM scoring functions is limited by the difficulty of including polarization and charge transfer effects into force fields [52], [55]–[58]. These effects typically appear in binding of small molecules to enzymes and metalloproteins, or when binding proceeds through a chemical reaction, such as proton transfer [59]. A quantum-mechanical description of the binding process naturally incorporates these effects, therefore, they can overcome these limitations. Unfortunately, a full *ab initio* simulation of the ligand-target complex is definitely beyond the capabilities of existing classical computers [52]. Furthermore, docking and screening require extensive sampling and evaluation of the score function, which would in turn require a significant number of *ab initio* calculations. These computational limitations have discouraged CADD practitioners to incorporate QM calculations in vHTS screening [59].

Despite the computational burden of QM calculation for CADD, different schemes have been devised to balance the accuracy requirements with the computation cost in the estimation of binding affinities. This task has been tackled mostly with Quantum Mechanics / Molecular Mechanics (QM/MM) approaches. In this case, the system is separated in two regions, one treated with a standard quantum chemistry approach (The QM region) and another region treated with MM. This approach is ideally complemented with a description of the solvent, either implicit, explicit or a combination of both [52]. The incorporation of these approaches to vHTS has been attempted via QM-based scoring functions that incorporate QM calculations for some parts of the system [60], as well as through simplified models of the binding [61]; however, these strategies have not been adopted by the CADD community mainly because of its larger computational cost and limited accuracy, the latter being influenced by the choice of methods and basis sets employed. As a result, the use of QM/MM approaches have been mostly prescribed to the lead optimization phase, where the set of molecules under study is significantly smaller. In the next section we describe in more detail the different approaches for the calculation of binding energies that incorporate QM calculations.

An alternative approach to vHTS for hit search is *de novo* design of ligands [62], [63]. These methods apply a strategy to build a completely new compound that can bind to the protein, generally by ligand growing or ligand linking methods. In the first approach, a known ligand is docked onto the binding pocket and additional groups are added or replaced on the initial structure to improve binding. In the second approach, a group of ligands is simultaneously docked onto the binding site and subsequently linked to generate a candidate. In both cases, a good scoring function coming from an *ab initio* simulation could improve the accuracy of the design; however, the biggest challenge for the *de novo* drug design is the assessment of synthesizability [5]: since the structures obtained in this process are new and do not appear in databases of synthesized compounds, it is important to establish a priori if it is worth devoting a significant amount of resources and expertise in synthesizing them. Current computational approaches to synthesizability employ automatic retrosynthetic studies [5]. Some of these tools could benefit from better QM approaches for simulating reaction paths and predicting the kinetics of chemical reactions, a space where quantum computers promise significant advances [64]. Similarly, state of the art machine learning techniques, such as generative adversarial networks [65], are offering new paths towards *de novo* design of small molecules [66], [67].

C. Lead discovery and optimization

Once hit compounds have been identified, they enter an optimization phase to produce a smaller set of better candidates, called *leads*. The set of leads undergoes further optimization in a process that iterates between CADD development and *in vitro* and animal experiments [5]. The purpose of this second phase of screening is to optimize the drug-like properties of the hit compounds, which includes not only the biological activity, but also the ADMET profile and other pharmacokinetic properties. The general assumption behind this process is that small changes on the chemical structure will produce incremental changes of the drug-like properties; therefore, the optimization involves the synthesis of the drug-candidates along with testing of their biological activities accompanied by CADD. In this stage, QSAR models for smaller datasets play a major role in the optimization, allowing for quickly judging whether certain modifications improve drug-likeness or not, especially when no target information is available.

On the other hand, structure-based approaches focus on estimating the binding affinity of the the candidate molecules, whose accurate and efficient calculation is considered the ultimate goal of CADD [59]. The binding affinity estimation needs to be very accurate because a difference of 6kJ/mol in this quantity translates to a change of one order of magnitude in the binding equilibrium constant, which finally determines the relative concentrations of the free-target and inhibited target. Unfortunately, most of the established methodologies for computing binding affinities are based on MM approximations that lack the description of polarization, charge transfer, and other quantum mechanical effects, among other phenomena crucial to achieve such accuracy. These limitations call for the use of quantum mechanical (QM) approaches, that naturally incorporate such effects. In practice, however, the use of QM calculations in drug discovery is not widespread due to the approximations that need to be employed to avoid the inherent computational cost of exact QM simulations on classical computers. Approximate QM methodologies can introduce errors that are even larger than those obtained with MM methods depending on the level of theory chosen and are usually restricted to only a small part of the ligand-protein complex due to computational limitations; consequently, the

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1147/JRD.2018.2888987, IBM Journal of Research and Development

development of efficient algorithms for accurately simulating quantum systems can have a significant impact on CADD.

Current CADD methods to compute binding affinities using QM methods can be classified as single-structure approaches, end-point approaches or full free-energy estimations [52]. In single-structured approaches, the energies are computed for a fixed ligand-target complex, with a structure obtained from experimental data or from a geometry optimization with MM. This approach avoids sampling multiple conformations with a QM energy function, but is very inaccurate due to the differences between the potential energy functions obtained with MM and QM approaches. This methodology can render errors of even 80 kJ/mol, even after relaxation of the geometries using the QM calculation. Most of the studies using a single-structure methodology have employed QM/MM and fragmentation approaches, where the QM part is treated by semiempirical methods (SE), density functional theory (DFT), or Hartree-Fock calculations with small basis sets. The accuracy observed in many applications can vary widely depending on the particular choice of QM method [52]. Fragmentation methods such as pairwise-additive (PA) fragmentation and molecular fragmentation with conjugate caps (MFCC) [68] and fragment molecular orbital (FMO) [69] have been employed in full calculations of the protein-ligand complex, with modest results. In all these methods, the total energy is approximated as a sum of the energy of individual fragments plus a correction term, which substantially the reduce the cost of the calculation. More recently, linear scaling methods have shown promise of increasing the accuracy, but are still far from end-point and full free-energy methodologies with force fields [52]. As a result, single-structure QM calculations have been employed mainly as a way to rationalize experimental binding affinities, rather than as a method for direct prediction.

In end-point approaches, the binding affinities are estimated as an average over several sampled structures, however, the expressions employed in the free energy calculation are empirical functions based on force field and empirical corrections to incorporate solvation effects. These methods have been upgraded by replacing the purely MM simulation with a QM/MM approach [52], [70]. To mitigate the computational cost of the application of QM/MM in drug discovery, the QM region is generally restricted to the ligand and occasionally some of the relevant residues in the binding pocket [52]. The QM approximations used in end-point approaches have been restricted to semi-empirical methods, density functional theory, or Hartree-Fock calculations.

The strict thermodynamic calculation of the binding free energies is much more challenging since it requires a free-energy simulation (FES) with extensive sampling of each of the intermediate molecular structures in the formation of the the receptor-ligand complex. In principle, the FES must describe the transition between the separate ligand and receptor structures to the receptor-ligand complex, but a more common approach is to compute relative binding energies by studying the transition between receptor-ligand complexes with different ligands. One strategy employed to avoid computing the QM/MM energy of all the sampled structures is to perform the entire calculation using MM and then to estimate a QM/MM correction using an average of the difference between the QM/MM and MM calculation for a few geometry configurations [71]–[73]. Another approach is the free energy perturbation (FEP) technique, where the first ligand is mutated by small structural changes into a second ligand to mimic the thermodynamic ligand-interconversion cycle, with averages taken over many samples [74]. While most of the FEP calculations are attempted with regular force fields, better accuracies could be achieved employing QM/MM techniques combined with implicit and explicit solvation models.

By far, FES approaches have proven to be the most accurate when applied in combination with QM methods of relatively high accuracy and sufficient sampling [52]. Benchmark calculations that compare the results of typical choices of QM methods such as DFT and semiempirical approaches have shown a clear improvement going from semiempirical, via DFT, to wavefunction based methods such as CCSD(T) [52]. Clearly, the major obstacle for approaches that combine FES and QM calculation for CADD is the daunting task of performing wavefunction calculations for all the structures sampled in FES. Moreover, properly accounting for solvation effects requires the combination of both explicit and implicit solvation techniques, which adds more complexity to the already challenging simulation of the receptor-ligand complex. In this scenario, the development of quantum algorithms that can tackle quantum simulation efficiently and accurately can contribute to extending the applicability of FES strategies to drug discovery, as quantum computers achieve size and precision that enables the simulation of large molecules.

II. Quantum computing

Most digital devices use *bits* as the building blocks for information processing. Each bit expresses a discrete, "classical" state of 0 or 1. Devices that perform computation by manipulating bits are referred to as *classical computers*. Quantum computers manipulate quantum states of matter for performing computation. A standard choice for constructing those quantum states is to combine two-level quantum systems called *qubits*. By manipulating the qubit states and taking advantage of uniquely quantum mechanical phenomena such as superposition and entanglement, quantum computers can perform computational tasks in ways that are beyond what is possible on their classical counterparts. A predefined way of manipulating quantum states to solve a computational problem is referred to as a *quantum algorithm*. In many cases, by analyzing the number of steps that quantum algorithms take it can be proved that they outperform classical algorithms for specific problems with reduced number of steps required. This capability is known as *quantum speedup*. Well-known examples of quantum speedup include Shor's algorithm for factorization, Grover search, and simulating quantum systems.

On the experimental side, a wide variety of physical systems have been explored as candidates for quantum com-

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1147/JRD.2018.2888987, IBM Journal of Research and Development

puters. Some of the hardware platforms, such as ion traps and superconducting qubits, have been scaling up rather rapidly in recent years towards the threshold regime beyond which it becomes intractable to simulate these physical systems with a classical computer. Recent results in both theory [75] and experiments [76] have pointed to how the so called *noisy intermediate-scale quantum* (NISQ) devices[1] [21] with moderate numbers of qubits can in principle produce quantum states whose measurement outcomes follow distributions that are justifiably hard to sample from on a classical computer.

Although quantum devices have scaled up rapidly in the past years, we must stress that the current quantum devices are still susceptible to noise and error due to environment interactions. This is the major obstacle that needs to be overcome along the path to scalable, fault-tolerant quantum computing (FTQC) devices. The good news is that a comprehensive and rigorous theory of quantum error correction has been developed over the past two decades which paves a plausible path to fault-tolerance. On the other hand, implementing most of the error correcting codes on NISQ hardware would incur a cost too prohibitive for the device to do anything useful other than error correction itself, leaving scalability and fault-tolerance a much longer term prospect than what is technologically feasible in the near future.

The presence of noise and error does pose significant challenges for quantum computing, however, it does not *a priori* exclude the possibility of achieving beyond-classical performance using NISQ devices. Noise rates and gate errors on NISQ machines have undergone a steady streak of improvement over the years. In addition, new ideas especially suited for taking advantage of NISQ devices [21] have provided interesting possibilities. The recently emerging paradigm of hybrid quantum-classical computation using variational quantum circuits has received wide interest [77]. In this paradigm, quantum computers are used as co-processors in tandem with classical computers for accomplishing simulation tasks (See Figure 2). One of the main appeals of this paradigm is that it combines the advantages of both quantum and classical computers. The quantum computer handles only the state preparation and measurement while the classical computer is tasked with optimizing the parameters with which the quantum computer uses for the state preparation.

In this section we focus on two large (and rapidly expanding) categories of quantum algorithms which are relevant to drug discovery, namely quantum algorithms for simulating molecular electronic structure in computational chemistry, and quantum-enhanced machine learning.

[1]NISQ devices are quantum computers that implement logical operations using physical qubits and, therefore, they have limited gate depth. Current NISQ computers are projected to have in the order of tens to hundreds of qubits and being able to execute circuits with depths in the order of thousands of two-qubit gates. While NISQ devices will not be able to implement error-correction, they are still expected to provide computational advantages over classical supercomputers for certain problems.

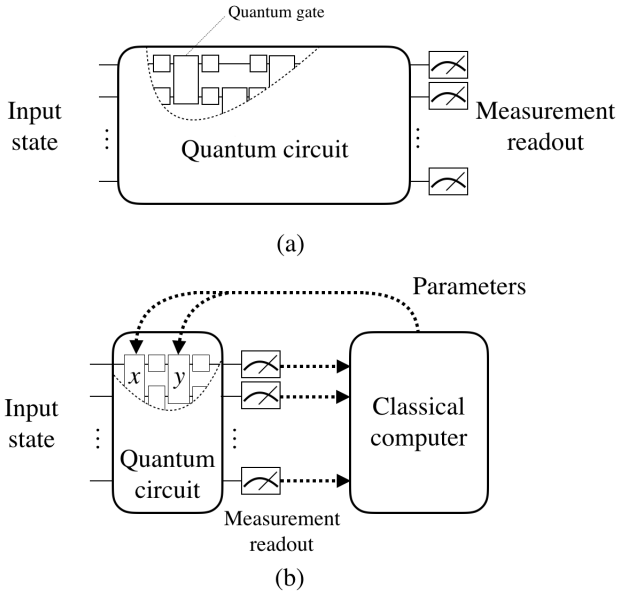

Fig. 2. Schematic diagrams of two quantum computing paradigms. (a) Standard gate model quantum computing. In the circuit diagram each horizontal wire represents a qubit and time flows from left to right. The quantum computation starts from some initial input state and the quantum computer applies a sequence of operations (called *quantum gates*) to generate a final state where each qubit is measured. The sequence of quantum gates is called a *quantum circuit*. (b) Variational quantum computation. Here the quantum circuit contains gates which have variable parameters (illustrated as x and y in the diagram). Each time the quantum circuit terminates, the classical computer gleans the measurement outcome and proposes new values of the parameters for the quantum circuit.

A. Quantum chemistry on quantum computers

Early results in the study of quantum algorithms have led to proposed quantum speedups for specific computational problems. Among these, simulation of quantum systems [9], [78]–[80] has been widely recognized as a canonical application for quantum computing. The quantum chemistry community has long observed that quantum mechanics is hard to simulate exactly on a classical computer [81]. The computational cost of exactly treating many-electron problems grows exponentially as the number of electrons increases. Feynman suggested the idea [78], [79], later formalized by Lloyd [80], that computers which obey quantum mechanics can efficiently simulating quantum systems. Broadly speaking, quantum simulation problems can be divided into static and dynamic problems. The former concerns computing the eigenvalues of the Hamiltonian and the latter concerns the time evolution of the wavefunction under a quantum Hamiltonian. Both tasks are difficult on a classical computer because the dimension of the quantum state space grows exponentially with respect to the size of the physical system. Unlike classical bits, n qubits can be in a superposition of 2^n states[2]. These n qubits may be in a highly entangled

[2]For example, three classical bits can only be in one of the eight possible states namely 000, 001, 010, $\cdots$, 111. However, three qubits can be in a state which is a superposition of all eight states: $a_{000}|000\rangle + a_{001}|001\rangle + \cdots + a_{111}|111\rangle$, corresponding to a unit vector $(a_{000}, a_{001}, \cdots, a_{111})$ of $2^3 = 8$ dimensions.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1147/JRD.2018.2888987, IBM Journal of Research and Development

state which can be efficiently generated and manipulated on a quantum computer but hard to treat on a classical computer. In the classical case, despite clever ideas for approximation [82], [83], in worst case scenarios one still needs to keep track of all of the 2^n amplitudes, which quickly becomes prohibitively expensive. For $n > 50$ storing the wavefunction of the qubits already requires some of the largest supercomputers in the world [84]–[86]. In contrast, by operating on the qubits in specific ways, quantum computers can efficiently simulate the dynamics of many quantum systems [80], and such ability also translates to the ability to efficiently obtain energy eigenvalues exactly [87], [88].

The unique features of quantum simulation have deep consequences for quantum chemistry. When doing quantum chemical calculations on classical computers, one would almost strictly avoid maintaining the explicit wavefunction of the physical system or propagating the full wavefunction unitarily under some quantum Hamiltonian due to the prohibitive costs of either method, while on a quantum computer state preparation and time evolution can often be done efficiently [80]. This distinction makes for rather different design patterns in quantum algorithms versus their classical counterparts. Building on previous results for quantum simulation [80], [87], [88], it was shown that using the ability to generate quantum states with sufficiently large overlap with the ground state and the ability to efficiently time evolve a state under a molecular Hamiltonian, one could obtain the ground state energy of the molecular Hamiltonian with accuracy comparable to full configuration interaction (FCI) (corresponding to exact diagonalization) [9]. While FCI suffers from exponential growth in computational cost, the cost of the quantum algorithm in [9] only scales polynomially with respect to the system size - an exponential improvement over classical algorithms.

Since the first introduction of quantum computing to quantum chemical applications [9], there has been substantial activity exploring quantum chemistry simulation in a variety of settings [89]–[93]. On the experimental side, simple demonstrations of quantum chemistry calculations for a diverse set of physical systems have been realized [94]–[98]. On the theoretical and numerical side, an extensive line of research has succeeded in reducing the resources required for quantum chemical simulations on a quantum computer [99]–[105]. Researchers have sought to map quantum simulation algorithms into concrete circuits and estimate the numbers of qubits and gates required for typical quantum chemistry calculations [64], [100]–[102]. There are also improvements to these quantum algorithms which are motivated by insights from quantum chemistry [103]–[105].

The quantum algorithms discussed so far assume that scalable, fault-tolerant quantum computers are available. This is a rather long-term prospect due to the extensive overhead required by quantum error correction [106], [107]. However, in the past few years, a hybrid quantum-classical scheme refered to as the *variational quantum eigensolver* (VQE) [77], [108] has received much attention due to its prospects for being deployed on impending NISQ devices. Unlike algorithms requiring long coherence time which is only possible on a fault-tolerant quantum computer, the basic idea of VQE is to use the quantum computer for a short state preparation step and use a classical computer to control the parameters with which the quantum computer prepares the state (Figure 2b). The measurement results coming from the quantum computer are often otherwise hard to compute on a classical computer. As a heuristic quantum algorithm for finding the ground state energy of a Hamiltonian, VQE operates by tuning the parameters of the quantum circuit to minimize the energy expectation of the output state with respect to the Hamiltonian. In recent years there have been a number of VQE experiments [98], [108]–[112] performed on a variety of physical platforms, some of which have already delivered results of chemical accuracy within a given one particle basis set for small systems such as berylium hydride.

An important aspect of quantum computing on NISQ devices is the incorporation of error mitigation techniques. Without error correction, the incoherent errors accumulated during the execution of the circuits translate into inaccurate expectation values measured for NISQ algorithms such as VQE. Recent proposals have focused on methods to estimate such errors and remove them from the final outcome of NISQ calculations. Unlike error-correction techniques, which are general techniques to indefinitely extend the coherence time of the quantum computer, error mitigation focuses specifically on compensating errors in expectation values and therefore do not require a large overhead on quantum resources. Most of the proposals have focused on methodologies to enhance the accuracy of expectation values by extrapolating from results with varying degrees of noise [113]–[115]. Another proposal suggests the possibility of estimating first order noise contributions to expectation values from calculations where only one qubit in the register is error-corrected [116]. More work is required to implement these proposals in combination with algorithms such as VQE, which will be crucial for achieving useful computation on NISQ devices.

One caveat with respect to VQE schemes, as well as variational quantum algorithms in general, is that training the quantum circuit from a random set of initial parameter guesses may lead to vanishingly small local gradient, making the optimization task challenging as the number of qubits grows [117]. This calls for good initial parameter guesses that are already somewhat "close" to the optimal solution. Fortunately for quantum chemical problems there are well motivated ansatz constructions that can serve as excellent initial starting points. In fact, the basic idea underlying some of the recent VQE proposals [108], [118] is to start from some initial approximation of the wavefunction that can be efficiently computed classically (Hartree-Fock [108] or fermionic Gaussian state [119]), prepare a quantum state on a quantum computer which represents such initial approximation, and variationally optimize the quantum circuit applied to the initial quantum state to yield a more refined quantum state that is otherwise intractable to realize classically (unitary coupled cluster [108] or fermionic non-Gaussian state

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1147/JRD.2018.2888987, IBM Journal of Research and Development

[119]). Such refinement can in principle provide results that are superior to what is feasible on a classical computer, as it corresponds to exploring a set of states for which there is not known efficient representation on a classical computer.

B. Quantum machine learning

In recent years there has been a rapidly expanding area of research seeking quantum techniques for enhancing machine learning methods. Although the full extent to which quantum computers can provide advantages on machine learning is far from known, there are a few heuristic arguments to support the belief that such advantages may exist:

1) Quantum computers can generate quantum states which give rise to probability distributions that are justifiably hard to sample from classically [75], [76]. Because of this ability to generate statistical patterns that are hard to generate classically, one hopes that quantum computers may also be able to recognize patterns in data that are hard to recognize classically [12].
2) For a physical system of n qubits, the space in which the quantum state of the n-qubit system dwells has dimension 2^n. Such exponential size may allow for an exponentially more compact encoding of classical information. For instance a quantum state of merely 30 qubits can represents a unit vector of length $2^{30} = 1,073,741,824$. In some cases, processing the 30-qubit state for machine learning purposes may be more advantageous than treating a vector of more than a billion entries on classical hardwares.
3) Many (classical) machine learning algorithms involve a large amount of linear algebraic operations, while quantum computers are known to provide speedups in problems related to some of the most elementary linear algebraic operations such as Fourier transforms [120, Ch. 5], vector inner products [121], matrix eigenvalues and eigenvectors [87], and solving linear systems of equations [122]. The quantum techniques therein can be used as a toolkit for building quantum machine algorithms.

The above intuitions underpin many of the existing quantum algorithms for machine learning tasks. Here we consider both supervised and unsupervised learning, which are common tools used in the early drug discovery processes. We also consider quantum machine learning techniques on both noisy intermediate scale quantum (NISQ) devices, and fault tolerant quantum computing (FTQC) devices. In the following discussions we highlight some representative results in each category and collect a more comprehensive list in Table I.

Supervised learning on NISQ devices. Some proposals in this category focus on taking advantage of the exponential size of the quantum state space (as alluded to in Argument 2 listed above) by encoding classical data into quantum states whose sizes scale logarithmically as the dimension of each data point [124], [125], [127]. It is further highlighted that one could train a quantum circuit with exponentially fewer parameters than what is possible with classical neural networks for classification tasks [124]. A small-scale demonstration of quantum classification has been realized recently on a quantum computer by IBM Corp. [125] for synthetic data. Similar ideas have also been implemented by Rigetti Computing, Inc. [128] using image recognition as an example, as well as by Xanadu Quantum Technologies, Inc. on a continuous variable quantum system [129]. In addition to using quantum circuits directly as classifiers, it has also been proposed to use a quantum computer to estimate hard-to-evaluate kernel functions for support vector machines [125], [126].

Supervised learning on FTQC devices. There are a variety of proposals in this category which take advantage of the linear algebraic routines developed in the quantum algorithm literature. Using techniques derived from Grover search, one can yield quantum speedup in training perceptrons [148] and restricted Boltzmann machines [142] on an FTQC devices. A similar speedup has also been shown for Bayesian inference [150].

The quantum algorithm for solving linear systems [122] opened up a new avenue from which quantum speedup may be obtained. For example, there are quantum algorithms for least-squares regression [144] and support vector machines [146] which may yield an exponential quantum speedup in certain settings.

Unsupervised learning on NISQ devices. There are algorithms in this category based on both gate model and quantum annealing devices. In the gate model, quantum autoencoder models [130], [131] have been proposed for learning compression of quantum data, which can be helpful for reducing the dimension of parameter space in variational quantum algorithms. Drawing on the connection to MAX-CUT, a clustering algorithm has been implemented [140] on a superconducting quantum device produced by Rigetti Computing, Inc. Another proposal for unsupervised quantum machine learning which is analogous to a Boltzmann machine uses the inherent property of sampling from a quantum circuit [157], and it has been demonstrated [138] on an ion trap quantum computer produced by IonQ, Inc.

Quantum annealing devices are also useful for unsupervised learning because they enable efficient approximate sampling from the thermal distributions of transverse Ising systems. The power of being able to approximately perform such sampling can be appreciated by observing that in the special case where the transverse field is absent, the sampling task is equivalent to a general Boltzmann machine in the classical setting, which is challenging in the classical case. Quantum annealers as samplers have been applied in generative models such as Helmholtz machines [136], variational autoencoders [133], and learning probabilistic graphical models [139] using processors produced by D-Wave Systems, Inc.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1147/JRD.2018.2888987, IBM Journal of Research and Development

	Supervised	Unsupervised
NISQ	Variational quantum circuit classifier [123]–[129] Kernel-based quantum-classical classifier [125], [126] Quantum Boltzmann machine [134], [135] Quantum training of classical Boltzmann machine [137]	Quantum autoencoder [130]–[132] Hybrid quantum-classical variational autoencoder [133] Hybrid quantum-classical Helmholtz machine [136] Quantum circuit-based generative modeling [138] Learning probabilistic graphical models [139] Quantum generative adversarial networks [118] Hybrid quantum-classical clustering [140] Quantum Boltzmann machine [134], [135] Quantum training of classical Boltzmann machine [137]
FTQC	Quantum-enhanced classical Boltzmann machine [141] Quantum nearest-neighbor classification [142] Quantum least-squares regression [144] Quantum support vector machine [146] Quantum perceptron models [148], [149] Quantum Bayesian inference [150] Quantum-enhanced Bayesian deep learning [151]	Quantum k-means clustering [142] Quantum principal component analysis [143] Quantum generative adversarial networks [145] Quantum Hopfield network [147] Quantum-enhanced classical Boltzmann machine [141]

TABLE I

Examples of techniques for using quantum computers for machine learning tasks. Here we broadly divide the algorithms into four groups. NISQ is short for Noisy Intermediate Scale Quantum devices [21], which are devices that are available currently or in the near term. FTQC refers to Fault Tolerant Quantum Computer, which are scalable, error-corrected devices that may be realized in more distant future. The algorithms listed in the former compartment have either been demonstrated on a physical device, or are designed with explicit consideration for deployment in NISQ devices. Although the algorithms listed in the FTQC category are not necessarily designed with NISQ devices in mind, they do deliver provable, asymptotic quantum advantage when scalable fault-tolerant quantum computers are available. However, we note that such claims of quantum advantage are not without caveats [12], [152], [153] especially in light of recent works [154]–[156] showing classical algorithms with comparable performance.

Unsupervised learning on FTQC devices. Many unsupervised learning tasks require computation of distance measures and estimating matrix eigenvalues. Indeed, there are quantum algorithms in this category which take advantage of the ability of quantum computers to efficiently perform linear algebraic operations. For example, on a fault-tolerant quantum computer one has access to quantum algorithms that allow for quadratic speedup in computing the inner product between two vectors compared with classical brute-force methods, which translates to quantum algorithm for k-means clustering that is faster than classical algorithms in certain regimes [142]. The ability to represent classical information in an exponentially compact way (Argument 2 in previous discussion) combined with the ability to efficiently analyze the spectrum of a matrix on a quantum computer (Argument 3) leads to a method for performing principal component analysis [143] in a way that is exponentially faster than classical computers in certain cases.

In summary, the use of quantum techniques for machine learning is an emerging field consisting of a diverse set of ideas. A few promising methods have been proposed recently which can be implemented on NISQ devices. However, much more utility would be unlocked if fully-error-corrected quantum computers become a reality (Table I). Of course, there are also challenges and caveats associated with what these quantum machine learning algorithms can achieve. For quantum machine learning algorithms building on the HHL algorithm [122] for linear systems of equations or similar ideas, caveats have been discussed in the literature [12], [152]. These raise the important questions of how to load classical data onto a quantum state (the input problem) and what meaningful results one could extract efficiently from the output quantum state (the output problem). Perhaps the most noteworthy caveat comes from a series of recent papers [154]–[156] which have developed a set of classical algorithms that deliver asymptotic performance comparable to the respective quantum algorithms [143], [158]. These advances demonstrate the subtleties associated with claims of quantum advantage in machine learning, but also how classical machine learning can benefit from quantum computing as a field.

Because NISQ devices mitigate error to a rather limited extent compared with the fault-tolerant quantum computers envisioned with the current theories, one major issue with quantum machine learning on NISQ devices is robustness to error and noise. Another challenge is the limited number of qubits that are available compared with the sizes and dimensions of real-world data sets where classical machine learning techniques struggle. For quantum annealers both challenges are discussed in detail by Perdomo-Ortiz et al. [153], where main issues include limited connectivity between the physical qubits, intrinsic noise and bias in device parameters, and deviation of the output distribution from the true Boltzmann distribution; nevertheless, empirical evidence shows that the annealing devices can still generate gradients for tuning the training parameters in the right direction [136]. In the gate model, gate error is a major concern. Nonetheless there is also evidence that variational quantum circuits can learn in the presence of error [137]. Recent years have witnessed the steady improvement of quantum hardware, bringing more qubits that are less error prone, combined with the development of hybrid quantum-classical techniques that reduce high dimensional real-world data to

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1147/JRD.2018.2888987, IBM Journal of Research and Development

lower dimensional latent space where samples from quantum devices are drawn [133], [136], [153]. These advances warrant cautious optimism for NISQ devices, in the near future, finding application to real-world machine learning tasks with competitive performances compared with existing classical techniques.

III. Opportunities for quantum computing in drug discovery

The potential to efficiently deliver quantum chemical calculations with accuracy comparable to FCI methods and find solutions to optimization problems can impact several of the areas of CADD describe above. Here we outline a few potential use cases for quantum computing (Figure 1c).

In structure-based drug discovery (Figure 1b), an important part of the input concerns the structure of the target protein. Some progress has been made in the past decade on quantum techniques for protein folding based on the aminoacid sequence. In particular, the quantum computing community has considered two simple models: the Hydrophobic-Polar (HP) model and Miyazawa-Jernigan (MJ) model, both of which model the protein as a self-avoided walk on a lattice. Solutions in both quantum annealers [159]–[162] and gate-model [163] quantum devices[3] have been explored. Current capability of these methods are limited to proof-of-concept examples such as Chignolin and Trp-Cage, small peptides with less than 21 aminoacid residues. Significant venture capital has been invested in quantum computing for life sciences[4] and future work will further unveil the magnitude of quantum advantage for larger protein folding problems as quantum devices scale up.

For molecular docking, one of the prevalent methods is atomistic modeling, which relies on force-field simplifications whose parameters need to match with quantum mechanical calculations. With the advent of variational quantum eigensolvers (VQE) and the quantum phase estimation algorithm (PEA), the size of physical systems that can be treated with accurate *ab initio* quantum calculations will be greatly expanded as quantum devices scale up. This allows for force field constructions based on exact quantum calculations for molecular fragments that are larger than what can be handled using existing quantum chemical methods.

In *de novo* design one of the pressing issues is synthesizability of a drug candidate, which involves simulation of different reaction paths. Quantum computers offer an avenue to potentially tackle electronic structure problems in the strongly correlated regime using, which would allow to simulate transition states and thermodynamic properties to accuracies comparable to FCI methods. As a result, this can improve the effectiveness of *de novo* design.

One of the bottlenecks for vHTS is the efficiency and accuracy with which one can calculate the scoring function (Section I-B). Ideally, the scoring function should be directly based on binding affinity, which comes from *ab initio* quantum mechanical calculations, while in practice empirical approximations are used. Hence with the quantum subroutine boosted by quantum computers, one may evaluate the scoring functions more efficiently and accurately. This could be achieved using methods where different parts of the system are computed with different levels of approximations, such as QM/MM [164]. The ability of computing binding affinities will also have a major impact on the lead optimization phase of drug discovery and mechanism of action studies, where understanding and quantitatively predicting the interaction of a drug candidate with multiple biological targets provides clues into toxicity, pharmacokinetics and multitarget action.

For ligand-based drug discovery, QSAR models have, in many cases, incorporated quantum mechanical properties [165], [166]. Generally, the quality and accuracy of these properties significantly affect the quality and predictivity of the model. Most of these approaches use descriptors derived from density-functional theory (DFT) calculations, and quantum computation could serve as a more efficient and accurate alternative for those calculations.

Another major aspect of QSAR is statistical and machine learning models. For example in virtual screening a common technique for classification in chemical space is by using kernels which map molecular structures to high dimensional features (see for example the "graph kernel" that has been used in cheminformatics literature [167], [168]). Commonly evaluating the kernel function requires handling vectors of extremely high dimensions, making computational efficiency a major issue for deploying kernel based classification methods [169]. In chemoinformatics, kernels accounting for the similarity between molecules are usually calculated from fingerprints or descriptor vectors using either some standard functions (linear, polynomial, Gaussian) or other popular similarity measures such as Euclidean distance or Tanimoto coefficient [170]. Positive-definite kernels for complex systems (*e.g.* protein-ligand complexes) can be constructed using a set of simple rules, as practiced in a subfield called *kernel engineering*). Since its inception, the kernel-based support vector machine (SVM) approach has become one of the most popular methods for building classification and regression structure-activity models [171], [172]. This is an area where quantum machine learning can enter with new possibilities for kernel designs. As mentioned in Section II-B, there are proposals for quantum enhanced support vector machines with kernel functions that are otherwise hard to evaluate on classical computers [125], [126]. The basic intuitions underlying the potential of having such kernel functions evaluated by quantum computers include the exponential size of the Hilbert space that a quantum state affords (see Argument 1 in Section II-B), as well as the ability to evaluate vector inner products efficiently (Argument 3). The precise regimes in which quantum algorithms may have an advantage over existing kernels, however, remains an open question of great value for future developments.

Most quantum simulation developments for quantum chemistry have focused on estimating molecular Hamiltonian

[3]There are NISQ devices in both quantum annealing and gate model architecture. For FTQC, there are constructions for gate model while for quantum annealing it is presently unclear.

[4]See for example ProteinQure, Inc. https://www.proteinqure.com/

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1147/JRD.2018.2888987, IBM Journal of Research and Development

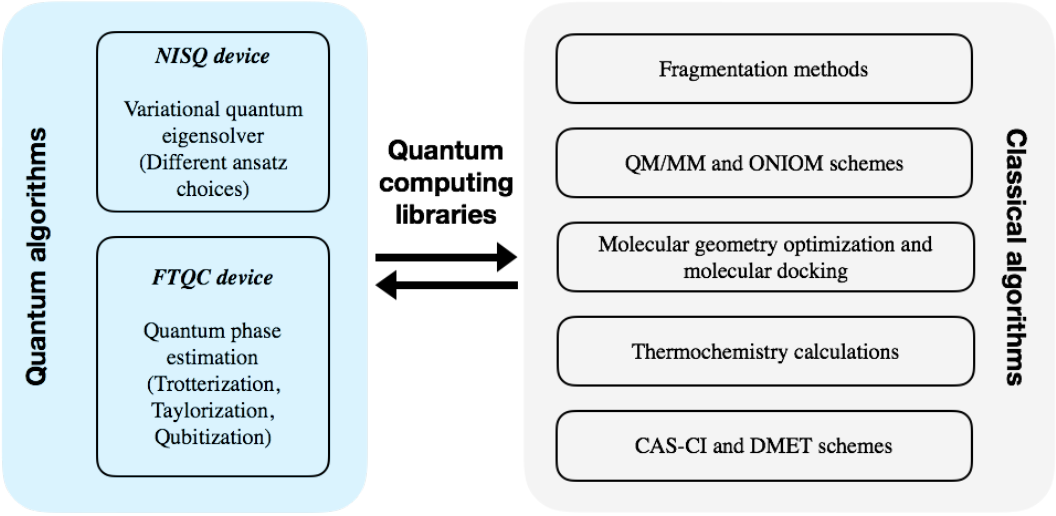

Fig. 3. As quantum devices improve to support the implementation of quantum algorithms for electronic structure calculations, the CADD community will take advantage of quantum computing by integrating quantum algorithms for computing molecular energies into existing classical quantum chemistry pipelines. The methods listed in the right can be employed for structure prediction, calculation of binding affinities, calculation of molecular descriptors for QSAR and to study reactions paths for *de novo* drug design (Figure 1b).

spectra and preparing eigenstates, with the exception of early proposals for the calculation of molecular properties [173], [174]. The main approach for this calculation on FTQC devices is quantum phase estimation, while approaches for NISQ devices are mainly based on the VQE algorithm. Quantum phase estimation approaches require the ability to implement a quantum operation which has eigenvalues that are a known function of the eigenvalues of the target Hamiltonian. The first methods used the dynamical evolution under the target Hamiltonian, that was implemented using techniques such as *Trotterization* [9], [91], [93], [99]–[101], [104], [105] and approximate decomposition into a sum of unitaries using a Taylor expansion (*Taylorization*) [92], [175]. More recently, an alternative approach called *qubitization* [176]–[178], where the evolution employs a quantum walk operator, has been proposed with near-optimal asymptotic performance. Current estimates using the qubitization technique indicate that calculations on systems of a thousand spin orbitals might require on the order of a million physical qubits with error rates of one in ten thousands using state of the art techniques for error correction, as described in [178]. Although further work is required to tighten these estimates and further optimize the use of quantum resources, these results indicate that electronic structure calculations are one of the first potential applications of FTQC devices.

While FTQC methods are guaranteed to deliver results of the same quality as FCI, the quality of the VQE calculations performed on NISQ devices depends on the ansatz employed. Generally, VQE is performed for the second quantized chemistry Hamiltonian, where the number of spin-orbitals corresponds to the number of physical qubits required in the calculation. For practical purposes, the VQE ansatz should not be efficiently realizable on a classical computer and preferably should have a depth that scales at most linearly with the size of the system. Fortunately, new ansatzes with linear depth have been recently proposed [85], [118], [179]. More work is required to implement and assess the accuracy and cost of these ansatzes for relevant molecular benchmark sets. This should include the study of thermochemical properties and interaction energies [180], [181]. Also, investigation of strongly correlated systems such as those involved in bond-breaking, excited states, and ground-state energy of transition metal complexes [182] could have direct applications in the pharmaceutical industry. These assessments must be complemented with comparisons against state of the art classical methods in electronic structure, such as density matrix renormalization group (DMRG) approaches [183], [184], density-matrix embedding theory (DMET) [185] and new developments in FCI calculations [186], [187].

Once available, energy solvers based on quantum computing algorithms could be integrated into hybrid schemes such as QM/MM, or in general, into methods where different parts of the system are treated using different levels of approximations, such as complete active space (CAS) calculations and ONIOM schemes [188] (See Figure 3). This integration requires the development of appropriate software interfaces between the quantum programs and computational chemistry programs. In this scheme, the quantum computer would act as a dedicated co-processor for exclusively accelerating parts of the calculation that require higher accuracy, connecting to classical computing hardware through software intermediate libraries. A first example of such libraries is OpenFermion [189], which incorporates tools for manipulating quantum chemistry operators, such as expressions in second quantization, and translate them into circuits and operators that can be implemented on a gate-based quantum computer. Along the same lines, software libraries for integrating quantum algorithms with existing machine learning libraries have

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1147/JRD.2018.2888987, IBM Journal of Research and Development

started to emerge [190]. More work in this direction will be essential in the development of useful quantum machine learning tools for CADD. Recently, the CADD community has been moving towards consolidating existing tools into python-based software platforms and towards the use of cloud computing services [191], which will facilitate their integration with existing quantum computing libraries.

As quantum libraries for controlling quantum computers start to emerge [192]–[194], more developments are required to cover the wide range of calculations common in computational chemistry, such as molecular geometry optimization, calculation of molecular properties, molecular spectra, tools for electronic density analysis, among others. While some of these calculations could be performed fully on a quantum processor [173] on an FTQC devices, an easier route in the NISQ era is to take advantage of the software interfaces described above to combine the quantum electronic wavefunction solver with classical routines for geometry optimization, calculation of vibrational modes, partition functions for thermodynamics, among others (See Figure 3). Many of these calculations require the estimation of derivatives of the electronic energy with respect to molecular coordinates and expectation values for the calculation of response properties. Consequently, developing efficient methods for measuring these properties on quantum computers is an immediate goal in the field of quantum computing for chemistry with direct implications in drug discovery.

As quantum hardware becomes more powerful, we expect quantum algorithms for chemistry and machine learning to be progressively integrated into CADD. While FTQC devices are not expected to be available within the next decade, NISQ devices would be more likely commercialized in the next 2-4 years [195]. The size of the problems that will be solved on these computers will be linked to their specifications (number of qubits and coherence time). For instance, quantum devices with qubit count N being a few hundreds of qubits and $O(N)$ coherence time would be able to perform quantum simulation on molecular systems with the same number of spin orbitals using VQE. Using techniques such as active space approach, we could study molecules of the size of typical drug candidates. This type of calculation could be useful in the parameterization of force-fields, in synthesizability and bio-catalysis studies and in the generation of QSAR descriptors. Calculation of binding energies using QM/MM techniques will likely require in the order of a few thousand qubits and will take advantage of the integration with classical tools, as described in Figure 3.

We are entering a new era of quantum computing where quantum hardware currently available already allow for rapid prototyping of quantum algorithms. As a result, the field is open to early explorations of how quantum devices can be used for concrete application settings. Drug discovery is a unique area in the sense that it benefits from advances in both quantum chemistry and machine learning, making it one of the first areas that are likely to adopt quantum computing into its pipelines. This perspective is an invitation to both the quantum computing and the drug discovery communities to bridge the technical gap needed to fully materialize the potential of quantum computing for drug discovery.

IV. Acknowledgements

We thank Laura Castellanos and Song Ping for insightful discussions about drug design and Tim Hirzel, Peter Johnson, Semion Saikin and Sukin Sim for their helpful comments on the manuscript. YC and AA-G acknowledge National Science Foundation Grant # CHE-1655187. JR and AA-G acknowledge the support from the Army Research Office under Award No. W911NF-15-1-0256. AA-G also acknowledges support from the Vannevar Bush Faculty Fellowship program sponsored by the Basic Research Office of the Assistant Secretary of Defense for Research and Engineering, under Award No. ONR 00014-16-1-2008.

References

[1] Joseph A DiMasi, Ronald W Hansen, and Henry G Grabowski. The price of innovation: new estimates of drug development costs. *J. Health Econ.*, 22(2):151–185, 2003.

[2] Thomas J Dougherty and Michael J Pucci. *Antibiotic discovery and development*. Springer Science & Business Media, 2011.

[3] *Estimated based on the FDA approvals from 2008 to 2017* Asher Mullard. 2017 FDA drug approvals, 2018.

[4] Vincent Zoete, Aurélien Grosdidier, and Olivier Michielin. Docking, virtual high throughput screening and in silico fragment-based drug design. *J. Cell. Mol. Med.*, 13(2):238–248, 2009.

[5] Gregory Sliwoski, Sandeepkumar Kothiwale, Jens Meiler, and Edward W Lowe. Computational methods in drug discovery. *Pharmacol. Rev.*, 66(1):334–395, 2014.

[6] John H Van Drie. Computer-aided drug design: the next 20 years. *J. Comput.-Aided Mol. Des.*, 21(10-11):591–601, 2007.

[7] Thomas Weymuth and Markus Reiher. Inverse quantum chemistry: Concepts and strategies for rational compound design. *Int. J. Quantum Chem.*, 114(13):823–837, 2014.

[8] Alán Aspuru-Guzik, Roland Lindh, and Markus Reiher. The matter simulation (r)evolution. *ACS Cent. Sci.*, 4(2):144–152, 2018.

[9] Alán Aspuru-Guzik, Anthony Dutoi, Peter Love, and Martin Head-Gordon. Simulated Quantum Computation of Molecular Energies. *Science*, 309(5741):1704–1707, 2005.

[10] Emily Conover. Google moves toward quantum supremacy with 72-qubit computer. *Sci. News*, 193(6), 2018.

[11] Chad Rigetti. The rigetti 128-qubit chip and what it means for quantum. *Medium*, Aug 2018. URL:https://medium.com/rigetti/the-rigetti-128-qubit-chip-and-what-it-means-for-quantum-df757d1b71ea.

[12] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum Machine Learning. *Nature*, 549(7671):195–202, 2016.

[13] J. P. Hughes, S. S. Rees, S. B. Kalindjian, and K. L. Philpott. Principles of early drug discovery. *Br. J. Pharmacol.*, 162(6):1239–1249, 2011.

[14] Kristian Stromgaard, Povl Krogsgaard-Larsen, and Ulf Madsen, editors. *Introduction to Drug Design and Discovery*. Taylor and Francis Group, Boca Raton, FL, 5th edition, 2017.

[15] Mario Compiani and Emidio Capriotti. Computational and theoretical methods for protein folding. *Biochemistry (Mosc.)*, 52(48):8601–8624, 2013.

[16] Yang Zhang. Progress and challenges in protein structure prediction. *Curr. Opin. Struct. Biol.*, 18(3):342–348, 2008.

[17] Subha Kalyaanamoorthy and Yi-Ping Phoebe Chen. Structure-based drug design to augment hit discovery. *Drug Discovery Today*, 16(17-18):831–839, 2011.

[18] Amy C Anderson. The process of structure-based drug design. *Chemistry & biology*, 10(9):787–797, 2003.

[19] Alexander Golbraikh, Xiang Simon Wang, Hao Zhu, and Alexander Tropsha. Predictive qsar modeling: methods and applications in drug discovery and chemical risk assessment. In *Handbook of computational chemistry*, pages 1309–1342. Springer, 2012.

[20] Jürgen Bajorath. Computer-aided drug discovery. *F1000Research*, 4, 2015.

[21] John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, Aug 2018.

[22] Brett Lomenick, Richard W Olsen, and Jing Huang. Identification of direct protein targets of small molecules. *ACS Chem. Biol.*, 6(1):34–46, 2010.

[23] Lyle Burdine and Thomas Kodadek. Target identification in chemical genetics: the (often) missing link. *Chemistry & biology*, 11(5):593–597, 2004.

[24] Janet NY Chan, Corey Nislow, and Andrew Emili. Recent advances and method development for drug target identification. *Trends Pharmacol. Sci.*, 31(2):82–88, 2010.

[25] Monica Schenone, Vlado Dančík, Bridget K Wagner, and Paul A Clemons. Target identification and mechanism of action in chemical biology and drug discovery. *Nat. Chem. Biol.*, 9(4):232, 2013.

[26] Xiaofeng S Zheng, Ting-Fung Chan, and Heather H Zhou. Genetic and genomic approaches to identify and study the targets of bioactive small molecules. *Chemistry & biology*, 11(5):609–618, 2004.

[27] Paul Czodrowski, Jan M Kriegl, Stefan Scheuerer, and Thomas Fox. Computational approaches to predict drug metabolism. *Expert opinion on drug metabolism & toxicology*, 5(1):15–27, 2009.

[28] Yongliang Yang, S. James Adelstein, and Amin I. Kassis. Target discovery from data mining approaches. *Drug Discovery Today*, 17(SUPPL.):S16–S23, 2012.

[29] Theodora Katsila, Georgios A Spyroulias, George P Patrinos, and Minos-Timotheos Matsoukas. Computational approaches in target identification and drug discovery. *Comput. Struct. Biotechnol. J.*, 14:177–184, 2016.

[30] Nidhi, Meir Glick, John W Davies, and Jeremy L Jenkins. Prediction of biological targets for compounds using multiple-category bayesian models trained on chemogenomics databases. *J. Chem. Inf. Model.*, 46(3):1124–1133, 2006.

[31] Jessica Lynn Grey and David H Thompson. Challenges and opportunities for new protein crystallization strategies in structure-based drug design. *Expert Opin. Drug Discovery*, 5(11):1039–1045, 2010.

[32] Frances C Bernstein, Thomas F Koetzle, Graheme JB Williams, Edgar F Meyer, Michael D Brice, John R Rodgers, Olga Kennard, Takehiko Shimanouchi, and Mitsuo Tasumi. The protein data bank: a computer-based archival file for macromolecular structures. *J. Mol. Biol.*, 112(3):535–542, 1977.

[33] Philip Bradley, Kira MS Misura, and David Baker. Toward high-resolution de novo structure prediction for small proteins. *Science*, 309(5742):1868–1871, 2005.

[34] R Laurie, T Alasdair, and Richard M Jackson. Methods for the prediction of protein-ligand binding sites for structure-based drug design and virtual ligand screening. *Curr. Protein Pept. Sci.*, 7(5):395–406, 2006.

[35] Stefan Henrich, Outi MH Salo-Ahen, Bingding Huang, Friedrich F Rippmann, Gabriele Cruciani, and Rebecca C Wade. Computational approaches to identifying and characterizing protein binding sites for ligand design. *J. Mol. Recognit.*, 23(2):209–219, 2010.

[36] Chayan Acharya, Andrew Coop, James E Polli, and Alexander D MacKerell. Recent advances in ligand-based drug design: relevance and utility of the conformationally sampled pharmacophore approach. *Curr. Comput.-Aided Drug Des.*, 7(1):10–22, 2011.

[37] Yovani Marrero-Ponce, Oscar Martínez Santiago, Yoan Martínez López, Stephen J Barigye, and Francisco Torrens. Derivatives in discrete mathematics: a novel graph-theoretical invariant for generating new 2/3d molecular descriptors. i. theory and qspr application. *J. Comput.-Aided Mol. Des.*, 26(11):1229–1246, 2012.

[38] Jens Auer and Jürgen Bajorath. Molecular similarity concepts and search calculations. In *Bioinformatics*, pages 327–347. Springer, 2008.

[39] Jürgen Bajorath. Integration of virtual and high-throughput screening. *Nat. Rev. Drug Discovery*, 1(11):882, 2002.

[40] Sean Ekins, Jordi Mestres, and Bernard Testa. In silico pharmacology for drug discovery: methods for virtual ligand screening and profiling. *Br. J. Pharmacol.*, 152(1):9–20, 2007.

[41] Shuxing Zhang. Computer-aided drug discovery and development. In *Drug Design and Discovery*, pages 23–38. Springer, 2011.

[42] Mingli Xiang, Yu Cao, Wenjie Fan, Lijuan Chen, and Yirong Mo. Computer-aided drug design: lead discovery and optimization. *Combinatorial chemistry & high throughput screening*, 15(4):328–337, 2012.

[43] Lu Zhang, Jianjun Tan, Dan Han, and Hao Zhu. From machine learning to deep learning: progress in machine intelligence for rational drug discovery. *Drug Discovery Today*, 2017.

[44] Lucy J Colwell. Statistical and machine learning approaches to predicting protein-ligand interactions. *Curr. Opin. Struct. Biol.*, 49:123–128, 2018.

[45] Angélica Nakagawa Lima, Eric Allison Philot, Gustavo Henrique Goulart Trossini, Luis Paulo Barbour Scott, Vinícius Gonçalves Maltarollo, and Kathia Maria Honorio. Use of machine learning approaches for novel drug discovery. *Expert Opin. Drug Discovery*, 11(3):225–239, 2016.

[46] Hongming Chen, Ola Engkvist, Yinhai Wang, Marcus Olivecrona, and Thomas Blaschke. The rise of deep learning in drug discovery. *Drug Discovery Today*, 2018.

[47] Mariya Popova, Olexandr Isayev, and Alexander Tropsha. Deep reinforcement learning for de novo drug design. *Science Advances*, 4(7):eaap7885, jul 2018.

[48] Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Cent. Sci.*, 2016.

[49] Thomas Luechtefeld, Dan Marsh, Craig Rowlands, and Thomas Hartung. Machine learning of toxicological big data enables read-across structure activity relationships (rasar) outperforming animal test reproducibility. *Toxicol. Sci.*, page kfy152, 2018.

[50] Inbal Halperin, Buyong Ma, Haim Wolfson, and Ruth Nussinov. Principles of docking: An overview of search algorithms and a guide to scoring functions. *Proteins: Struct., Funct., Bioinf.*, 47(4):409–443, 2002.

[51] Michael L Connolly. Analytical molecular surface calculation. *J. Appl. Crystallogr.*, 16(5):548–558, 1983.

[52] Ulf Ryde and Par Soderhjelm. Ligand-binding affinity estimates supported by quantum-mechanical methods. *Chem. Rev.*, 116(9):5520–5566, 2016.

[53] Thomas A Halgren. Merck molecular force field. i. basis, form, scope, parameterization, and performance of mmff94. *J. Comput. Chem.*, 17(5-6):490–519, 1996.

[54] Noel M OBoyle, John W Liebeschuetz, and Jason C Cole. Testing assumptions and hypotheses for rescoring success in protein- ligand docking. *J. Chem. Inf. Model.*, 49(8):1871–1878, 2009.

[55] Arjan van der Vaart and Kenneth M Merz. The role of polarization and charge transfer in the solvation of biomolecules. *J. Am. Chem. Soc.*, 121(39):9182–9190, 1999.

[56] Arjan van der Vaart and Kenneth M Merz Jr. Charge transfer in biologically important molecules: comparison of high-level ab initio and semiempirical methods. *Int. J. Quantum Chem.*, 77(1):27–43, 2000.

[57] Richard A Friesner. Modeling polarization in proteins and protein–ligand complexes: Methods and preliminary results. *Advances in protein chemistry*, 72:79–104, 2005.

[58] Richard A Bryce. Physics-based scoring of protein–ligand interactions: explicit polarizability, quantum mechanics and free energies. *Future medicinal chemistry*, 3(6):683–698, 2011.

[59] Kaushik Raha, Martin B Peters, Bing Wang, Ning Yu, Andrew M Wollacott, Lance M Westerhoff, and Kenneth M Merz Jr. The role of quantum mechanics in structure-based drug design. *Drug Discovery Today*, 12(17-18):725–731, 2007.

[60] Li Rao, Igor Ying Zhang, Wenping Guo, Li Feng, Eric Meggers, and Xin Xu. Nonfitting protein–ligand interaction scoring function based on first-principles theoretical chemistry methods: Development and application on kinase inhibitors. *J. Comput. Chem.*, 34(19):1636–1646, 2013.

[61] Ting Zhou and Amedeo Caflisch. High-throughput virtual screening using quantum mechanical probes: Discovery of selective kinase inhibitors. *Chemmedchem*, 5(7):1007–1014, 2010.

[62] Markus Hartenfeller and Gisbert Schneider. De novo drug design. In *Chemoinformatics and computational chemical biology*, pages 299–323. Springer, 2010.

[63] Gisbert Schneider. De novo design–hop (p) ing against hope. *Drug Discovery Today: Technologies*, 10(4):e453–e460, 2013.

[64] Markus Reiher, Nathan Wiebe, Krysta M. Svore, Dave Wecker, and Matthias Troyer. Elucidating reaction mechanisms on quantum computers. *Proceedings of the National Academy of Sciences*, 2017.

[65] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

[66] Peter Sadowski, David Fooshee, Niranjan Subrahmanya, and Pierre Baldi. Synergies between quantum mechanics and machine learning in reaction prediction. *J. Chem. Inf. Model.*, 56(11):2125–2128, 2016.

[67] Evgeny Putin, Arip Asadulaev, Yan Ivanenkov, Vladimir Aladinskiy, Benjamin Sánchez-Lengeling, Alán Aspuru-Guzik, and Alex Zhavoronkov. Reinforced adversarial neural computer for de novo molecular design. *J. Chem. Inf. Model.*, 2018.

[68] Da W Zhang and JZH Zhang. Molecular fractionation with conjugate caps for full quantum mechanical calculation of protein–molecule interaction energy. *The Journal of chemical physics*, 119(7):3599–3605, 2003.

[69] Dmitri G Fedorov and Kazuo Kitaura. Extending the power of quantum chemistry to large systems with the fragment molecular orbital method. *J. Phys. Chem. A*, 111(30):6904–6914, 2007.

[70] Claudio Norberto Cavasotto, Natalia Sol Adler, and María Gabriela Aucar. Quantum chemical approaches in structure-based virtual screening and lead optimization. *Front. Chem.*, 6:188, 2018.

[71] V Luzhkov and A Warshel. Microscopic models for quantum mechanical calculations of chemical processes in solutions: Ld/ampac and scaas/ampac calculations of solvation energies. *J. Comput. Chem.*, 13(2):199–213, 1992.

[72] Thomas H Rod and Ulf Ryde. Quantum mechanical free energy barrier for an enzymatic reaction. *Phys. Rev. Lett.*, 94(13):138302, 2005.

[73] Fernanda Duarte, Beat A Amrein, David Blaha-Nelson, and Shina CL Kamerlin. Recent advances in qm/mm free energy calculations using reference potentials. *Biochimica et Biophysica Acta (BBA)-General Subjects*, 1850(5):954–965, 2015.

[74] Robert Abel, Lingle Wang, Edward D Harder, BJ Berne, and Richard A Friesner. Advancing drug discovery through enhanced free energy calculations. *Acc. Chem. Res.*, 50(7):1625–1632, 2017.

[75] Scott Aaronson and Lijie Chen. Complexity-theoretic foundations of quantum supremacy experiments. 2016. arXiv:1612.05903.

[76] Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J. Bremner, John M. Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, apr 2018.

[77] Jarrod R. McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.*, 18(2):1–20, 2016.

[78] Richard P. Feynman. Simulating physics with computers. *Int. J. Theor. Phys.*, 21(6-7):467–488, 1982.

[79] Richard P. Feynman. Quantum mechanical computers. *Foundations of Physics*, 16(6):507–531, Jun 1986.

[80] Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996.

[81] W. Kohn. Nobel Lecture: Electronic structure of matterwave functions and density functionals. *Rev. Mod. Phys.*, 71(5):1253–1266, 1999.

[82] Román Orús. A practical introduction to tensor networks: Matrix product states and projected entangled pair states. *Ann. Phys.*, 349:117–158, 2014.

[83] Igor L. Markov and Yaoyun Shi. Simulating Quantum Computation by Contracting Tensor Networks. *SIAM J. Comput.*, 38(3):963–981, 2008.

[84] Zhaoyun Chen, Qi Zhou, Cheng Xue, Xia Yang, Guangcan Guo, and Guoping Guo. 64-Qubit Quantum Circuit Simulation. *Science Bulletin*, 63(15):964–971, 2018.

[85] Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, and Hartmut Neven. Simulation of low-depth quantum circuits as complex undirected graphical models, 2017. arXiv:1712.05384.

[86] Edwin Pednault, John A. Gunnels, Giacomo Nannicini, Lior Horesh, Thomas Magerlein, Edgar Solomonik, and Robert Wisnieff. Breaking the 49-qubit barrier in the simulation of quantum circuits, 2017. arXiv:1710.05867.

[87] Daniel S Abrams and Seth Lloyd. Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors. *Phys. Rev. Lett.*, 83(24):5162, 1999.

[88] A. Yu. Kitaev. Quantum measurements and the abelian stabilizer problem, 1995. arXiv:quant-ph/9511026.

[89] I Kassal, S P Jordan, P J Love, M. Mohseni, and A. Aspuru-Guzik. Polynomial-time quantum algorithm for the simulation of chemical dynamics. *Proceedings of the National Academy of Sciences*, 105(48):18681–18686, 2008.

[90] Ryan Babbush, Peter J. Love, and Alán Aspuru-Guzik. Adiabatic Quantum Simulation of Quantum Chemistry. *Sci. Rep.*, 4(6603):6603, 2014.

[91] Ryan Babbush, Dominic W Berry, Yuval R Sanders, Ian D Kivlichan, Artur Scherer, Annie Y Wei, Peter J Love, and Alán Aspuru-Guzik. Exponentially more precise quantum simulation of fermions in the configuration interaction representation. *Quantum Science and Technology*, 3(1):015006, 2017.

[92] Ryan Babbush, Dominic W. Berry, Ian D. Kivlichan, Annie Y. Wei, Peter J. Love, and Alán Aspuru-Guzik. Exponentially more precise quantum simulation of fermions in second quantization. *New J. Phys.*, 18(3):1–12, 2016.

[93] Ian D Kivlichan, Nathan Wiebe, Ryan Babbush, and Alán Aspuru-Guzik. Bounding the costs of quantum simulation of many-body physics in real space. *J. Phys. A: Math. Theor.*, 50(30):305301, 2017.

[94] 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, A. Aspuru-Guzik, and A. G. White. Towards quantum chemistry on a quantum computer. *Nat. Chem.*, 2(2):106–111, 2010.

[95] Zhaokai Li, Man-Hong Yung, Hongwei Chen, Dawei Lu, James D Whitfield, Xinhua Peng, Alán Aspuru-Guzik, and Jiangfeng Du. Solving Quantum Ground-State Problems with Nuclear Magnetic Resonance. *Sci. Rep.*, 1(November):88, 2011.

[96] Ya Wang, Florian Dolde, Jacob Biamonte, Ryan Babbush, Ville Bergholm, Sen Yang, Ingmar Jakobi, Philipp Neumann, Aln Aspuru-Guzik, James D. Whitfield, and Jrg Wrachtrup. Quantum simulation of helium hydride cation in a solid-state spin register. *ACS Nano*, 9(8):7769–7774, 2015. PMID: 25905564.

[97] Jacob Mower, Nicholas C. Harris, Gregory R. Steinbrecher, Yoav Lahini, and Dirk Englund. High-fidelity quantum state evolution in imperfect photonic integrated circuits. *Phys. Rev. A*, 92:032322, Sep 2015.

[98] P. J J O'Malley, Ryan Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, Alán Aspuru-Guzik, and J. M. Martinis. Scalable quantum simulation of molecular energies. *Phys. Rev. X*, 6(3):1–13, 2016.

[99] James D. Whitfield, Jacob Biamonte, and Alán Aspuru-Guzik. Simulation of electronic structure Hamiltonians using quantum computers. *Mol. Phys.*, 109(5):735–750, 2011.

[100] Dave Wecker, Bela Bauer, Bryan K. Clark, Matthew B. Hastings, and Matthias Troyer. Gate-count estimates for performing quantum chemistry on small quantum computers. *Physical Review A - Atomic, Molecular, and Optical Physics*, 90(2):1–15, 2014.

[101] David Poulin, Matthew B. Hastings, Dave Wecker, Nathan Wiebe, Andrew C. Doherty, and Matthias Troyer. The Trotter Step Size Required for Accurate Quantum Simulation of Quantum Chemistry. *Quantum Information & Computation*, 15(5-6):13, 2014.

[102] Matthew B. Hastings, Dave Wecker, Bela Bauer, and Matthias Troyer. Improving Quantum Algorithms for Quantum Chemistry. *Quantum Information & Computation*, 15(1-2):1–21, 2014.

[103] Jonathan Romero, Ryan Babbush, Jarrod R McClean, Cornelius Hempel, Peter J Love, and Alán Aspuru-Guzik. Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz. *Quantum Science and Technology*, 4(1):014008, 2018.

[104] Ryan Babbush, Jarrod McClean, Dave Wecker, Alán Aspuru-Guzik, and Nathan Wiebe. Chemical basis of Trotter-Suzuki errors in quantum chemistry simulation. *Physical Review A - Atomic, Molecular, and Optical Physics*, 91(2):1–18, 2015.

[105] Jarrod R. McClean, Ryan Babbush, Peter J. Love, and Alán Aspuru-Guzik. Exploiting locality in quantum computation for quantum chemistry. *J. Phys. Chem. Lett.*, 5(24):4368–4380, 2014.

[106] Daniel Gottesman. Stabilizer codes and quantum error correction, 1997. arXiv:quant-ph/9705052.

[107] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Physical Review A - Atomic, Molecular, and Optical Physics*, 86(3), 2012.

[108] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O'Brien. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.*, 5(May):4213, 2014.

[109] Yangchao Shen, Xiang Zhang, Shuaining Zhang, Jing-Ning Zhang, Man-Hong Yung, and Kihwan Kim. Quantum implementation of the unitary coupled cluster for simulating molecular electronic structure. *Physical Review A*, 95(2), Feb 2017.

[110] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549(7671):242, 2017.

[111] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi. Computation of molecular spectra on a quantum processor with an error-resilient algorithm. *Phys. Rev. X*, 8:011021, 2018.

[112] Cornelius Hempel, Christine Maier, Jonathan Romero, Jarrod McClean, Thomas Monz, Heng Shen, Petar Jurcevic, Ben P. Lanyon, Peter Love, Ryan Babbush, Alán Aspuru-Guzik, Rainer Blatt, and Christian F. Roos. Quantum chemistry calculations on a trapped-ion quantum simulator. *Phys. Rev. X*, 8:031022, 2018.

[113] Kristan Temme, Sergey Bravyi, and Jay M Gambetta. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.*, 119(18):180509, 2017.

[114] Suguru Endo, Simon C. Benjamin, and Ying Li. Practical quantum error mitigation for near-future applications. *Physical Review X*, 8(3), Jul 2018.

[115] Abhinav Kandala, Kristan Temme, Antonio D. Corcoles, Antonio Mezzacapo, Jerry M. Chow, and Jay M. Gambetta. Extending the computational reach of a noisy superconducting quantum processor, 2018. arXiv:1805.04492.

[116] Matthew Otten and Stephen Gray. Accounting for errors in quantum algorithms via individual error reduction, 2018. arXiv:1804.06969.

[117] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes, 2018. arXiv:1803.11173.

[118] Pierre-Luc Dallaire-Demers and Nathan Killoran. Quantum generative adversarial networks. *arXiV:1804.09139*, pages 1–23, 2018.

[119] Pierre-Luc Dallaire-Demers, Jonathan Romero, Libor Veis, Sukin Sim, and Alán Aspuru-Guzik. Low-depth circuit ansatz for preparing correlated fermionic states on a quantum computer. 2018. arXiv:1801.01053.

[120] Michael A Nielsen and Isaac Chuang. *Quantum Computation and Quantum Information*. 2000.

[121] Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Phys. Rev. Lett.*, 87(16):167902, 2001.

[122] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.*, 103(15):1–8, 2009.

[123] Edward Farhi and Hartmut Neven. Classification with quantum neural networks on near term processors. 2018. arXiv:1802.06002.

[124] Maria Schuld, Alex Bocharov, Krysta Svore, and Nathan Wiebe. Circuit-centric quantum classifiers. 2018. arXiv:1804.00633.

[125] Vojtech Havlicek, Antonio D Córcoles, Kristan Temme, Aram W Harrow, Jerry M Chow, and Jay M Gambetta. Supervised learning with quantum enhanced feature spaces. 2018. arXiv:1804.11326.

[126] Maria Schuld and Nathan Killoran. Quantum machine learning in feature Hilbert spaces. *arXiv:1803.07128*, 2018.

[127] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii. Quantum circuit learning. *Phys. Rev. A*, 98:032309, Sep 2018.

[128] C. M. Wilson, J. S. Otterbach, N. Tezak, R. S. Smith, G. E. Crooks, and M. P. da Silva. Quantum Kitchen Sinks: An algorithm for machine learning on near-term quantum computers. *arXiv:1806.08321*, pages 1–8, 2018.

[129] Nathan Killoran, Thomas R Bromley, Juan Miguel Arrazola, Maria Schuld, Nicolás Quesada, and Seth Lloyd. Continuous-variable quantum neural networks. 2018. arXiv:1806.06871.

[130] Jonathan Romero, Jonathan Olson, and Alan Aspuru-guzik. Quantum autoencoders for efficient compression of quantum data. *arXiv:1612.02806*, pages 1–8, 2016.

[131] Kwok Ho Wan, Oscar Dahlsten, Hlér Kristjánsson, Robert Gardner, and MS Kim. Quantum generalisation of feedforward neural networks. *npj Quantum Inf.*, 3(1):36, 2017.

[132] Lucas Lamata, Unai Alvarez-Rodriguez, José David Martín-Guerrero, Mikel Sanz, and Enrique Solano. Quantum autoencoders via quantum adders with genetic algorithms. *Quantum Science and Technology*, 4(1):014007, 2018.

[133] Amir Khoshaman, Walter Vinci, Brandon Denis, Evgeny Andriyash, and Mohammad H Amin. Quantum variational autoencoder. 2018. arXiv:1802.05779.

[134] Mohammad H. Amin, Evgeny Andriyash, Jason Rolfe, Bohdan Kulchytskyy, Roger Melko, David Goldberg, John Holland, M I Jordan, T M Mitchell, David C Parkes, Michael P Wellman, Jrgen Schmidhuber, Adi L Tarca, Vincent J Carey, Xue-wen Chen, Roberto Romero, and Sorin Drăghici. Quantum Boltzmann Machine. *PLoS Comput. Biol.*, 3(6245):267–272, 2015.

[135] Mária Kieferová and Nathan Wiebe. Tomography and generative training with quantum boltzmann machines. *Phys. Rev. A*, 96:062327, Dec 2017.

[136] Marcello Benedetti, John Realpe-Gómez, and Alejandro Perdomo-Ortiz. Quantum-assisted helmholtz machines: a quantum–classical deep learning framework for industrial datasets in near-term devices. *Quantum Science and Technology*, 3(3):034007, 2018.

[137] Guillaume Verdon, Michael Broughton, and Jacob Biamonte. A quantum algorithm to train neural networks using low-depth circuits. 2017. arXiv:1712.05304.

[138] Marcello Benedetti, Delfina Garcia-Pintos, Yunseong Nam, and Alejandro Perdomo-Ortiz. A generative modeling approach for benchmarking and training shallow quantum circuits. 2018. arXiv:1801.07686.

[139] Marcello Benedetti, John Realpe-Gómez, Rupak Biswas, and Alejandro Perdomo-Ortiz. Quantum-assisted learning of hardware-embedded probabilistic graphical models. *Phys. Rev. X*, 7(4):52–54, 2017.

[140] JS Otterbach, R Manenti, N Alidoust, A Bestwick, M Block, B Bloom, S Caldwell, N Didier, E Schuyler Fried, S Hong, et al. Unsupervised machine learning on a hybrid quantum computer. 2017. arXiv:1712.05771.

[141] Nathan Wiebe, Ashish Kapoor, and Krysta M. Svore. Quantum Deep Learning. *e-print arXiv: 1412.3489*, 2(3):1–34, 2014.

[142] Nathan Wiebe, Ashish Kapoor, and Krysta M. Svore. Quantum algorithms for nearest-neighbor methods for supervised and unsupervised learning. *arXiv:1401.2142v2*, pages 1–31, 2014.

[143] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum principal component analysis. *Nat. Phys.*, 10(9):631–633, 2014.

[144] Nathan Wiebe, Daniel Braun, and Seth Lloyd. Quantum Data Fitting. pages 1–6, 2012. arXiv:1204.5242.

[145] Seth Lloyd and Christian Weedbrook. Quantum generative adversarial learning. 2018. arXiv:1804.09139.

[146] Patrick Rebentrost, Masoud Mohseni, and Seth Lloyd. Quantum support vector machine for big data classification. *Phys. Rev. Lett.*, 113(3):1–5, 2014.

[147] Patrick Rebentrost, Thomas R Bromley, Christian Weedbrook, and Seth Lloyd. A quantum recurrent neural network. 2017. arXiv:1710.03599.

[148] Nathan Wiebe, Ashish Kapoor, and Krysta M Svore. Quantum Perceptron Models. *arXiv*, 0(x):4, 2016.

[149] Yudong Cao, Gian Giacomo Guerreschi, and Alán Aspuru-Guzik. Quantum neuron: an elementary building block for machine learning on quantum computers. 2017. arXiv:1711.11240.

[150] Guang Hao Low, Theodore J. Yoder, and Isaac L. Chuang. Quantum inference on Bayesian networks. *Physical Review A - Atomic, Molecular, and Optical Physics*, 89(6), 2014.

[151] Zhikuan Zhao, Alejandro Pozas-Kerstjens, Patrick Rebentrost, and Peter Wittek. Bayesian deep learning on a quantum computer, 2018. arXiv:1806.11463.

[152] Scott Aaronson. Quantum Machine Learning Algorithms : Read the Fine Print. *Nat. Phys.*, 11:1–5, 2015.

[153] Alejandro Perdomo-Ortiz, Marcello Benedetti, John Realpe-Gómez, and Rupak Biswas. Opportunities and challenges for quantum-assisted machine learning in near-term quantum computers. *Quantum Science and Technology*, 3(3):030502, 2018.

[154] Ewin Tang. A quantum-inspired classical algorithm for recommendation systems, 2018. arXiv:1807.04271.

[155] Ewin Tang. Quantum-inspired classical algorithms for principal component analysis and supervised clustering, 2018. arXiv:1811.00414.

[156] András Gilyén, Seth Lloyd, and Ewin Tang. Quantum-inspired low-rank stochastic regression with logarithmic dependence on the dimension, 2018. arXiv:1811.04909.

[157] Song Cheng, Jing Chen, and Lei Wang. Information perspective to probabilistic modeling: Boltzmann machines versus born machines. *Entropy*, 20(8):583, 2018.

[158] Iordanis Kerenidis and Anupam Prakash. Quantum recommendation systems, 2016. arXiv:1603.08675.

[159] Alejandro Perdomo, Colin Truncik, Ivan Tubert-Brohman, Geordie Rose, and Aln Aspuru-Guzik. Construction of model Hamiltonians for adiabatic quantum computation and its application to finding low-energy conformations of lattice protein models. *Physical Review A - Atomic, Molecular, and Optical Physics*, 78(1), 2008.

[160] Alejandro Perdomo-Ortiz, Neil Dickson, Marshall Drew-Brook, Geordie Rose, and Aln Aspuru-Guzik. Finding low-energy conformations of lattice protein models by quantum annealing. *Sci. Rep.*, 2(April):1–7, 2012.

[161] Ryan Babbush, Alejandro Perdomo-Ortiz, Bryan O'Gorman, William Macready, and Alan Aspuru-Guzik. Construction of Energy Functions for Lattice Heteropolymer Models: Efficient Encodings for Constraint Satisfaction Programming and Quantum Annealing. *Adv. Chem. Phys.*, 155:201–244, 2014.

[162] Tomáš Babej, Mark Fingerhuth, et al. Coarse-grained lattice protein folding on a quantum annealer. 2018. arXiv:1811.00713.

[163] Mark Fingerhuth, Tomáš Babej, et al. A quantum alternating operator ansatz with hard and soft constraints for lattice protein folding. 2018. arXiv:1810.13411.

[164] Olayide A Arodola and Mahmoud ES Soliman. Quantum mechanics implementation in drug-design workflows: does it really help? *Drug Design, Development and Therapy*, 11:2551, 2017.

[165] Ahmed El Kerdawy, Stefan Güssregen, Hans Matter, Matthias Hennemann, and Timothy Clark. Quantum mechanics-based properties for 3D-QSAR. *J. Chem. Inf. Model.*, 53(6):1486–1502, 2013.

[166] Steve Dixon, Kenneth M. Merz, Giorgio Lauri, and James C. Ianni. QMQSAR: Utilization of a semiempirical probe potential in a field-based QSAR method. *J. Comput. Chem.*, 26(1):23–34, 2005.

[167] Liva Ralaivola, Sanjay J. Swamidass, Hiroto Saigo, and Pierre Baldi. Graph kernels for chemical informatics. *Neural Networks*, 18(8):1093–1110, 2005.

[168] Pierre Mahé, Nobuhisa Ueda, Tatsuya Akutsu, Jean-Luc Perret, and Jean-Philippe Vert. Graph Kernels for Molecular StructureActivity Relationship Analysis with Support Vector Machines. *J. Chem. Inf. Model.*, 45(4):939–951, 2005.

[169] Te-Ming Huang, Vojislav Kecman, and Ivica Kopriva. *Kernel based algorithms for mining huge data sets*, volume 1. Springer, 2006.

[170] B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, 2002.

[171] Nianyi Chen. *Support vector machine in chemistry*. World Scientific, 2004.

[172] Ovidiu Ivanciuc. Applications of support vector machines in chemistry. *Reviews in computational chemistry*, 23:291, 2007.

[173] Ivan Kassal and Alán Aspuru-Guzik. Quantum algorithm for molecular properties and geometry optimization. *The Journal of chemical physics*, 131(22):224102, 2009.

[174] Daniel A Lidar and Haobin Wang. Calculating the thermal rate constant with exponential speedup on a quantum computer. *Physical Review E*, 59(2):2429, 1999.

[175] Ryan Babbush, Nathan Wiebe, Jarrod McClean, James McClain, Hartmut Neven, and Garnet Kin-Lic Chan. Low-depth quantum simulation of materials. *Phys. Rev. X*, 8(1):011044, 2018.

[176] Guang Hao Low and Isaac L Chuang. Hamiltonian simulation by qubitization. 2016. arXiv:1610.06546.

[177] Dominic W Berry, Mária Kieferová, Artur Scherer, Yuval R Sanders, Guang Hao Low, Nathan Wiebe, Craig Gidney, and Ryan Babbush. Improved techniques for preparing eigenstates of fermionic hamiltonians. *NPJ Quantum Inf.*, 4(1):22, 2018.

[178] Ryan Babbush, Craig Gidney, Dominic W Berry, Nathan Wiebe, Jarrod McClean, Alexandru Paler, Austin Fowler, and Hartmut Neven. Encoding electronic spectra in quantum circuits with linear t complexity. 2018. arXiv:1805.03662.

[179] Ian D Kivlichan, Jarrod McClean, Nathan Wiebe, Craig Gidney, Alán Aspuru-Guzik, Garnet Kin-Lic Chan, and Ryan Babbush. Quantum simulation of electronic structure with linear depth and connectivity. *Phys. Rev. Lett.*, 120(11):110501, 2018.

[180] Russell D Johnson III. NIST 101. Computational Chemistry Comparison and Benchmark Database. Technical report, 2018. https://cccbdb.nist.gov/.

[181] Lars Goerigk and Stefan Grimme. Efficient and accurate double-hybrid-meta-gga density functionals evaluation with the extended gmtkn30 database for general main group thermochemistry, kinetics, and noncovalent interactions. *J. Chem. Theory Comput.*, 7(2):291–309, 2010.

[182] Ricardo A Mata and Martin A Suhm. Benchmarking quantum chemical methods: Are we heading in the right direction? *Angew. Chem. Int. Ed.*, 56(37):11011–11018, 2017.

[183] Garnet Kin-Lic Chan and Sandeep Sharma. The density matrix renormalization group in quantum chemistry. *Annu. Rev. Phys. Chem.*, 62:465–481, 2011.

[184] Örs Legeza, Thorsten Rohwedder, Reinhold Schneider, and Szilárd Szalay. Tensor product approximation (dmrg) and coupled cluster method in quantum chemistry. In *Many-Electron Approaches in Physics, Chemistry and Mathematics*, pages 53–76. Springer, 2014.

[185] Sebastian Wouters, Carlos A Jiménez-Hoyos, Qiming Sun, and Garnet K-L Chan. A practical guide to density matrix embedding theory in quantum chemistry. *J. Chem. Theory Comput.*, 12(6):2706–2719, 2016.

[186] George H Booth, Alex JW Thom, and Ali Alavi. Fermion Monte Carlo without fixed nodes: A game of life, death, and annihilation in Slater determinant space. *The Journal of chemical physics*, 131(5):054106, 2009.

[187] Deidre Cleland, George H Booth, and Ali Alavi. Communications: Survival of the fittest: Accelerating convergence in full configuration-interaction quantum Monte Carlo, 2010.

[188] Lung Wa Chung, WMC Sameera, Romain Ramozzi, Alister J Page, Miho Hatanaka, Galina P Petrova, Travis V Harris, Xin Li, Zhuofeng Ke, Fengyi Liu, et al. The ONIOM method and its applications. *Chem. Rev.*, 115(12):5678–5796, 2015.

[189] Jarrod R McClean, Ian D Kivlichan, Damian S Steiger, Yudong Cao, E Schuyler Fried, Craig Gidney, Thomas Häner, Vojtěch Havlíček, Zhang Jiang, Matthew Neeley, et al. Openfermion: the electronic structure package for quantum computers. 2017. arXiv:1710.07629.

[190] Ville Bergholm, Josh Izaac, Maria Schuld, Christian Gogolin, and Nathan Killoran. Pennylane: Automatic differentiation of hybrid quantum-classical computations. 2018. arXiv:1811.04968.

[191] Chanin Nantasenamat and Virapong Prachayasittikul. Maximizing computational tools for successful drug discovery, 2015.

[192] Robert S Smith, Michael J Curtis, and William J Zeng. A practical quantum instruction set architecture. 2016. arXiv:1608.03355.

[193] Andrew W Cross, Lev S Bishop, John A Smolin, and Jay M Gambetta. Open quantum assembly language. 2017. arXiv:1707.03429.

[194] Dave Wecker and Krysta M Svore. LIQUi$|\rangle$: A software design architecture and domain-specific language for quantum computing. 2014. arXiv:1402.4467.

[195] Masoud Mohseni, Peter Read, Hartmut Neven, Sergio Boixo, Vasil Denchev, Ryan Babbush, Austin Fowler, Vadim Smelyanskiy, and John Martinis. Commercialize quantum technologies in five years. *Nature News*, 543(7644):171, 2017.