

Advances of machine learning in molecular modeling and simulation

Mojtaba Haghaghatlari¹ and Johannes Hachmann^{1,2,3}



In this review, we highlight recent developments in the application of machine learning for molecular modeling and simulation. After giving a brief overview of the foundations, components, and workflow of a typical supervised learning approach for chemical problems, we showcase areas and state-of-the-art examples of their deployment. In this context, we discuss how machine learning relates to, supports, and augments more traditional physics-based approaches in computational research. We conclude by outlining challenges and future research directions that need to be addressed in order to make machine learning a mainstream chemical engineering tool.

Addresses

¹ Department of Chemical and Biological Engineering, University at Buffalo, The State University of New York, Buffalo, NY 14260, United States

² Computational and Data-Enabled Science and Engineering Graduate Program, University at Buffalo, The State University of New York, Buffalo, NY 14260, United States

³ New York State Center of Excellence in Materials Informatics, Buffalo, NY 14203, United States

Corresponding authors:

Haghaghatlari, Mojtaba (mojtabah@buffalo.edu),
Hachmann, Johannes (hachmann@buffalo.edu)

Current Opinion in Chemical Engineering 2019, **23**:xx-yy

This review comes from a themed issue on **Frontiers of chemical engineering: molecular modeling**

Edited by **Randall Snurr, Jim Pfaendtner** and **Veronique Van Speybroeck**

<https://doi.org/10.1016/j.coche.2019.02.009>

2211-3398/© 2019 Elsevier Ltd. All rights reserved.

Machine learning from a chemical perspective

Over the past few years, data science has started to offer a fresh perspective on tackling complex chemical questions, such as discovering and designing chemical systems with tailored property profiles, revealing intricate structure-property relationships (SPRs), and exploring the vastness of chemical space [1[•]]. Data-derived prediction models serve as surrogates for physics-based models that are at the heart of traditional modeling and simulation work. They are attractive, because they are usually

dramatically less demanding than physics-based models and can thus be deployed in studies of correspondingly larger scope and scale. If trained on experimental data, they are also not subject to the approximations made in physics-based models and may thus not exhibit the resulting discrepancies with respect to non-idealized experimental findings. Of course, data-derived models have their own intrinsic errors and limitations, which we will address in the course of this review.

Machine learning (ML) is a data mining technique and used to create data-derived models. It enables us to extract complex and often hidden correlations (and thus ideally insights, patterns, rules, and guidance) from given data sets and to encapsulate them in mathematical form. ML is commonly categorized into four types, i.e., supervised, semi-supervised, unsupervised, and reinforcement learning. The main difference between these types is in essence the amount of information (i.e., labeling, context) that is available for the target variable that serves as the ground truth for the training of an ML algorithm. While all ML types have found application in chemical research [2], supervised learning has so far been most commonly used, and this review will thus focus on it. The popularity of supervised learning may be due to its heuristic and intuitive approach to learning, which is similar to a scientist's way of gaining insights into SPR. A supervised prediction model can be thought of as a function $f: X \rightarrow Y$ that maps an input $x \in X$ to an output $y \in Y$, where x in this context is the feature representation of a chemical system and y its target property. If the variables x and y are continuous (numerical), then the mapping is a regression; if they are discrete (categorical), then it is a classification.

We can utilize a host of supervised ML algorithms to train and optimize model f to approximate the output value for a given input. Two popular algorithms that have been widely used are artificial neural networks (ANNs) and kernel methods. Both can be thought of as transforming the input x into a new feature (latent variable) space, in which it becomes linearly correlated with the output y [3]. The transformation itself is typically highly non-linear. A major advantage of the ANNs is their capacity to transform features sequentially through several layers, which is referred to as deep learning. Kernel methods, on the other hand, usually transform features in a one-step process using kernel functions. Unlike in ANNs, this process is predefined prior to the tuning of the model's parameters, and is thus less flexible to learn the best latent variable space. The advantage of kernel methods is their superior

performance in finding global solutions, even for small-size data sets where ANNs have deficits. The support vector machines and kernel ridge regression are two common examples of kernel-based algorithms.

The relationship between a molecular structure and its properties is deterministic, i.e., there exists an exact mapping from fundamental physics (i.e., the Schrödinger equation). This mapping is ultimately the foundation for traditional modeling and simulation techniques. The topology of ML models is generally very flexible (as, e.g., expressed in the universal approximation theorem for ANNs), so that they can learn and recover the underlying SPRs of a problem, even from simple chemical representations (assuming no significant loss of information within this representation).

We can consider a feature representation method as a function $g: M \rightarrow X$ that maps a basic chemical representation $m \in M$ to a feature input $x \in X$ (typically called a descriptor). The representation m may contain spatial or at least topological information that defines a molecule and is expressed, e.g., in atomic coordinates, simplified molecular-input line-entry system (SMILES) [4], international chemical identifier (InChI), or other formats.

A common feature space X is spanned by structural descriptors. Some ML approaches also utilize physical or (physico-)chemical properties as descriptors, such that g corresponds to a simulation or some other type of calculation (including those from *first principles*). As this approach builds physics into the feature space, it has a certain appeal and has gained corresponding popularity. However, it is important to point out that the computational cost of obtaining such descriptors (which include optimized geometries) may easily make this the bottleneck of an ML approach, in which case it will limit its utility as an efficient surrogate for the prediction of y . This issue has to be considered as part of a cost-benefit analysis.

Another class of descriptors is designed to capture the local environment of each atom in a molecule [5]. This approach considers a molecule as a graph with atom and bond (i.e., node and edge) features. Each atom can interact with all other atoms in its immediate vicinity, which results in an update of the corresponding local atomic features. Incidentally, this approach has its roots in both chemical and data sciences: In the context of molecular simulations, cutoff radii have long been used to exploit the short-ranged nature of intermolecular interactions. In data science, the idea of dynamic irregular graphs provides the underpinning of graph convolutional neural networks. The overlap of the two disciplines in this area has led to many methodological advances for the generation of descriptors. Results from a number of recent studies suggest that an ensemble of local features

(rather than a global representation), is able to provide a more robust solution to the challenges involved with variant graph size and the order of atoms in molecules [6,7].

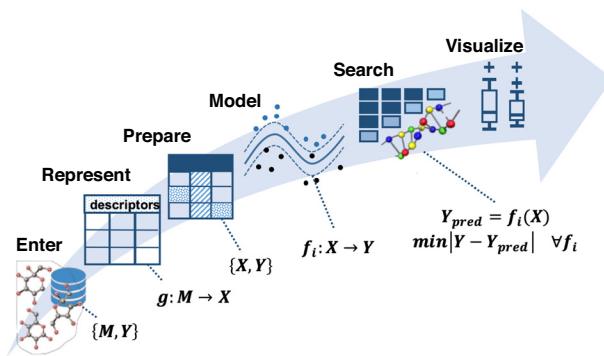
The descriptors discussed so far are essentially hand-crafted to explicitly expose certain structural, physical, or (physico-)chemical information x from m and provide a structured (i.e., tabular) feature representation. Alternatively, the feature generation g can also be merged into the prediction model f and both will be jointly optimized, e.g., through hidden layers (latent space) in deep learning. This class of descriptors is called learned features [8].

The overall ML workflow for chemical problems encompasses a number of steps as shown in Figure 1, including parsing, cleaning, and preprocessing a chemical data set $\{M, Y\}$, compiling an array of descriptors via g , as well as training, evaluation, and validation of the prediction model f .

Applications of machine learning in chemical research

In the following section, we summarize three application areas of ML in chemical research, with particular consideration of the inherent structure of the associated data sets, types of representation, and connections to traditional modeling. We limit the scope of our discussion to molecular systems, which still cover a broad range of use cases.

Figure 1



Current Opinion in Chemical Engineering

The major tasks and mathematical setup of a supervised machine learning workflow: For a given data set $\{M, Y\}$, in which for a number of molecules in basic chemical representation $m \in M$ the target property $y \in Y$ is given, we apply a feature representation method as a function $g: M \rightarrow X$ that maps M to a feature input space X . After cleaning and other preprocessing, we use $\{X, Y\}$ to formulate an ML model $f: X \rightarrow Y$ that maps the feature input space X to the output label space Y . The ML model is trained on the training subset of $\{X, Y\}$, and subsequently validated and optimized on its testing subset, so that it minimizes the prediction error for Y .

Discovery and design of new compounds

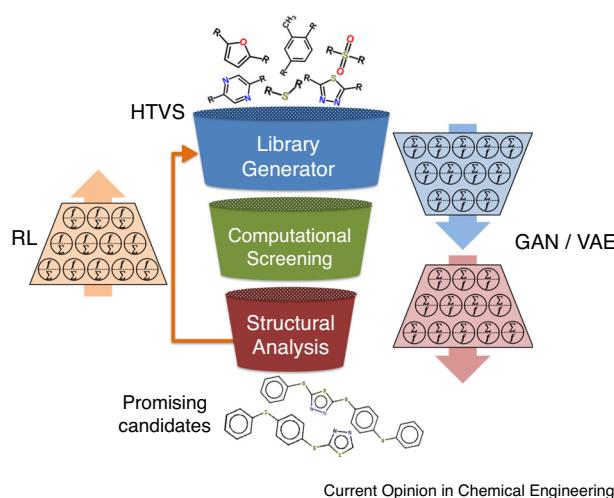
The application of ML for the exploration of chemical space and the creation of new compounds (ranging from small molecules to polymers and materials) can be divided into two distinct approaches: (i) discovery, i.e., ML is used to generate fast prediction models for properties of interest, with which large-scale surveys of chemical space can be conducted in order to identify compounds that exhibit desired property profiles; (ii) design, i.e., ML is used to develop a quantitative understanding of the SPRs of interest, which can be inverted to pursue the targeted, rational design (or inverse engineering) of compounds with particular properties. While the core activity, i.e., the ML of SPRs, is the same in both cases, its use follows different perspectives.

Discovery. The idea of employing data-derived prediction models instead of physics-based models (or experimentation) as a means to characterize candidates in the search for new molecules may be one of the earliest applications in chemistry, for which the use of ML was proposed. Traditional molecular modeling and simulations have been used for this purpose for many years. More recently, they have also been employed in the context of virtual high-throughput screening studies, in which they are tasked with assessing entire libraries of candidate compounds (see Figure 2 and, e.g., Refs. [9,10]). However, the computational footprint, in particular of *first-principles* approaches, is limiting both individual as well as large-

scale studies that seek to identify compounds with specifically targeted properties.

The application of data-derived prediction models enables us to dramatically accelerate the survey of chemical space, often by several orders of magnitude. A speed-up of that magnitude allows a corresponding increase in the scale and scope that is viable for screening efforts. (It is thus sometimes referred to as *hyperscreening*.) The candidate libraries are typically generated from a collection of moieties and patterns that are of interest in a given context [11]. The combination of such a set of building blocks leads to a molecular library for a particular domain in chemical space, i.e., the candidates belong to the same distribution [12]. A number of experimental or high-level computational training sets have been developed for specific classes of molecules [13,14]. Since these data sets focus on relatively similar compounds from the same distribution, the choice of representation and the ML model training are arguably less challenging compared to more universally applicable models. The extrapolative use of data-derived prediction models outside the domain for which they were trained has to be conducted with great care and caution, as they are least reliable here. This is a conceptual challenge, as screening studies are often interested in compounds with extreme properties that are likely at the margins of the distribution, where the predictions are least reliable. Iterative retraining of ML models allows us to shift the training data distribution into particular areas of interest, thus making them more robust for use in discovery.

Figure 2



Flowchart showing a computational funnel typical for high-throughput virtual screening (HTVS) studies. The neural network schemes on the left and right represent deep generative model architectures that can conceptually replace different elements of the screening funnel. Both generative adversarial networks (GANs) and variational autoencoders (VAEs) include two networks that revolutionize the conventional generation and analysis steps by probabilistic means. A deep reinforcement learning (RL) network can also be trained to bias the generation towards promising candidates.

A reasonably diverse collection of molecules can be found in the open-source QM9 data set originally extracted from the GDB-17 chemical universe of 166 billion organic molecules [15]. The QM9 data comprises computed geometries and properties for 134 000 molecules at density functional theory (DFT) level. Due to the diversity of molecular structures and broad range of calculated properties, QM9 plays an important role as a benchmark data set for new models and methods [16,17]. Its contribution to method developments can be compared to the MNIST data set in the hand-written character recognition community [18]. In contrast to, e.g., data sets from *first-principles* modeling, those from data-derived models have so far rarely been used for the generation of new reference data. Yet, they have played an important role in a number of methodological advances in the field. As a result, the reported accuracies for many of the recent ML prediction models surpass those of traditional molecular modeling and simulations [19].

Design. While discovery is still based on a traditional trial-and-error process — albeit one drastically accelerated by ML — the notion of a deliberate, *de novo* design of new compounds represents a different research paradigm. It addresses the problem that even rapid and efficient

hyperscreening studies can only scratch the surface of the practically infinite molecular space. Instead, the design paradigm seeks to utilize insights into the SPRs obtained from ML for the targeted creation of systems with specific properties. The understanding of how changes in the molecular structure (or a compound's features) lead to changes in the desired properties can be inverted to gain a property to structure mapping. The mathematical structure of a data-derived SPR prediction model (e.g., the dominant features, principal components, latent variables, or learned features) yields a foundation for inverse design. Models that are less easy to interpret can be projected onto surrogate models for which the extraction of guidelines is easier. A key challenge is to realize the simultaneous enhancement of different properties. The emerging design rules can be used to formulate individual compounds [20[•]], but also to identify high-value domains in chemical space, which can be enumerated in screening libraries (e.g., by sampling compounds similar to a lead compound). The latter approach is effectively interfacing the discovery and design perspectives and allows both physics-based and data-derived modeling studies to be more targeted.

Another approach that is very promising and has caused much excitement is the application of generative models (see Figure 2). For instance, Sanchez-Lengeling *et al.* have shown that a generative adversarial network (GAN) in tandem with reinforcement learning can outperform evolutionary algorithms in order to bias the generative process toward the extreme regions of a property distribution [21,22]. The use of GANs for molecular design and library generation is very recent and a number of concerns and challenges still need to be overcome in their development. Two of the principal challenges of GANs (and other generative approaches) are the rate at which invalid (i.e., chemically irrelevant or non-sensical) structures are generated, and their ability to produce topologically different molecules compared to the underlying training data [23,24]. Another example of generative models are variational autoencoders (VAEs) that learn the distribution of embedded space and thus enables tuning in that space [25]. Recurrent neural networks (RNNs) operate in a sequential manner similar to creating new molecules one atom at a time. One benefit of RNNs is their memory mechanism that allows them to remember the effects of previous sequences [26].

Creation of new modeling techniques

Instead of replacing physics-based with data-derived modeling entirely as outlined in the section 'Discovery and design of new compounds', ML can also be used to (i) *calibrate and correct* the results of physics-based models to account for some of their systematic errors; (ii) *complement* traditional modeling and simulation approaches (i.e., employ combinations of physics-based and data-derived

models); and (iii) facilitate the *development of new physics-based modeling techniques*.

The *calibration* approach allows us to improve the predictive performance of physics-based models and obtain high-quality results at the cost of lower-quality methods. It can also help bridge the gap between experiment and theory that results from the inherent approximations in the latter (see, e.g., Ref. [10]). Transfer learning is an ML design methodology that has been a particularly successful technique in this context [27,28]. In the *combination* approach, we only utilize ML for aspects for which no good physics-based models are available or where their use is impractical (e.g., because of insufficient accuracy, prohibitive cost, or other numerical issues). We thus retain as much of the physical foundations and robustness of traditional modeling as possible, while being pragmatic about the parts of a problem, where that is not possible (see, e.g., Refs. [29–31]).

The *development of entirely new modeling techniques* by means of ML has seen encouraging pioneering efforts, in particular for force fields (FFs) and DFT. The major driving force behind ML-generated FFs is the lack of generalizability in the classical FFs and the interatomic potentials that underpin them. This is an area where ML is apparently able to bridge the accuracy and versatility typically seen from quantum chemistry and the efficiency of molecular mechanics simulations. A recent line of research has focused on learning interatomic potentials from quantum chemical data sets [32]. There are two specific challenges involved in this application that make it distinct from prediction models for molecular properties. One is the need for a diverse sampling of non-equilibrium chemical conformations, as both ML and classical FFs perform poorly outside of their applicability domain. Access to a diverse collection of high-quality training samples is thus essential in creating ML FFs. For instance, Botu *et al.* have improved on previous work by diversifying their training data, e.g., by adding more atomic environments and applying clustering methods [33]. Smith *et al.* have pushed the normal mode sampling method to obtain single point energies for more than 20 million conformations generated for 58 000 small molecules [34[•]]. The results of these efforts were shown to be efficiently generalizable, even for the simulation of more complex phenomena. The second important challenge is to conserve the consistency between potential energies and forces as discussed by Chmiela *et al.* [35]. They provide a robust solution to this challenge by developing gradient-domain ML models (which reproduce global FFs by training in the force domain and incorporating both energies and forces) in an automated fashion, thus learning accurate ML FFs.

In the DFT context, ML is used to create new functionals for different terms in the electronic Hamiltonian. The

exact form of several functionals (e.g., the kinetic energy functional for interacting electrons or the exchange-correlation functional in the Kohn–Sham formalism) are unknown and otherwise approximated by physical reasoning. The ML-generated functionals allow DFT to avoid common failures, such as in accurately describing bond-breaking processes. Different ML functionals for specific classes of molecules, target properties, and electronic structure situations are being developed, as are fast methods that, e.g., learn energy functionals directly without having to solve the Kohn–Sham equations, thus making them a viable approach for *ab initio* molecular dynamics simulations [36].

Predictions of chemical reactions and catalyst systems

Research on chemical reactions is another field that has been benefiting from the advances in ML methodology. ML has been paving the way for a better understanding of chemical transformations with numerous real-world implications. SMILES are often the representation of choice for both the inputs (reactants and reagents) and outputs (products) of data-derived models for chemical reactions. These models are trained on known reactions to recognize structural patterns that may undergo bond-breaking or -formation in the course of a reaction or catalytic process [37*]. One particularly important data set for this application domain is the result of the US patent reaction extraction by Lowe [38].

The progress in predicting organic reactions and their products has been particularly noteworthy in recent years. Nam *et al.* have introduced sequence-to-sequence models to address the reaction prediction task similar to linguistic translation problems [39]. More recently, Schwaller *et al.* outperformed a similar approach in an end-to-end template-free model with a focus on the attention mechanism and a new tokenization strategy [40]. Coley *et al.* introduced a graph convolutional neural network approach with competitive performance. It was used for the prediction of reaction products as well as reactive sites of the reagents that are most likely to initiate a reaction [41]. One major contribution of the last two studies is the development of web applications to facilitate easy access to their models. These tools are available via the IBMRXN and ASKCOS websites, respectively [42,43].

A promising direction of ongoing work is the prediction of reaction pathways and mechanisms. All these efforts ultimately aim for a practical and more generalizable implementation of retrosynthetic analysis, which has been a grand challenge in organic chemistry for many years [44]. Insights regarding the synthetic feasibility of virtual compounds are also a key concern for the screening library generation and molecular design efforts discussed in the section ‘Discovery and design of new compounds’.

Outlook on future directions

Feature representations

As discussed in the section ‘Machine learning from a chemical perspective’, the descriptors of a given molecular system are an abstraction of its detailed nature (as well as a numerical representation). The choice of a suitable feature space is still our first and most effective means to infuse physics into ML models. There have been efforts to define criteria for the development of efficient descriptors [5], e.g., that they are (1) invariant to the symmetries of the underlying physics; (2) easy to interpret; (3) expressed in a direct and concise form to avoid redundancy and the curse of dimensionality; and (4) computationally efficient. However, developing molecular representations that adhere to all these criteria has been an exceedingly difficult task. More importantly, there is now agreement that ML approaches may require different types of descriptors to recover the entirety of SPRs of molecular systems. Further research into the creation of new descriptors (including fingerprint schemes) as well as the formulation of additional criteria will be necessary for the foreseeable future. The accessibility and flexibility of deep learning models can accelerate future developments via learned features and theory-informed models.

Machine learning for small data

While ML ideas became popular during the recent ‘big data’ wave (i.e., in chemistry with the emergence of large-scale screening result from high-level *first-principles* modeling), large data sets are in practice more often than not unavailable. In fact, problems for which data is (still) sparse tend to be of particular interest. As the data generation (both from experiment and modeling) is often a limiting factor, we will have to strive to reduce its cost or the number of data points needed to obtain ML models of a desired accuracy. It is thus essential to put an emphasis on developing ML methods that achieve better performance on small data sets. As mentioned in the section ‘Creation of new modeling techniques’, transfer learning is a promising approach in this context. We will also need to employ smart sampling methods and identify data points that are most important for the training of ML models. Active learning strategies offer a path towards this goal [45–47]. Many of these techniques are of general-purpose utility, but some will have to be tailored towards the specific problem settings of data-derived models for chemistry.

Software and tool development

The idea to utilize ML and other data mining techniques in the chemical domain is so recent that much of the basic infrastructure has not yet been developed, or is still in its early stages [1*]. The majority of tools and expertise tend to be technically involved, labor intensive, or otherwise unavailable to the community at large. Many researchers are now starting to pursue open-source software development projects to tackle this situation [48*]. However, the

lack of rigorous development guidelines remains a challenge that researchers from domain science need to overcome to make their efforts lasting and sustainable. The Molecular Sciences Software Institute (MoSSI) is one of the pioneers in establishing best practices and guidance for early-stage software developments in this field [49,50].

Conclusions

In this review, we discussed how ML can advance traditional modeling and simulation by (partially) replacing them (i.e., choosing data-derived over physics-based models or combining the two); calibrating, augmenting, or otherwise correcting their results; targeting studies and their objectives; and providing the means to effectively mine their results for a deeper understanding of hidden SPRs. Many ML models are still built on data provided by modeling and simulation — often as part of virtual high-throughput screening studies — and combining ML and traditional modeling infuses physics and robustness into the resulting data-derived prediction models. These and other emerging ML techniques have been enabling accelerated discovery and rational design in numerous areas of chemistry. Its early successes indicate that ML is bound to become a mainstream tool in chemical research. Yet, there is still much to (machine) learn on how to develop the full potential of ML in chemistry.

Conflict of interest statement

Nothing declared.

Acknowledgements

MH gratefully acknowledges support by Phase-I and Phase-II Software Fellowships (grant no. ACI-1547580-479590) of the National Science Foundation (NSF) Molecular Sciences Software Institute (grant no. ACI-1547580) at Virginia Tech. JH acknowledges support by the NSF CAREER program under grant no. OAC-1751161, the NSF Big Data Spokes program under grant no. IIS-1761990, and funding by the New York State Center of Excellence in Materials Informatics (grant no. CMI-1148092-8-75163).

References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest

1. Hachmann J, Windus TL, McLean JA, Allwardt V, Schrimpe-Rutledge AC, Afzal MAF, Haghhighatlari M: *Framing the Role of Big Data and Modern Data Science in Chemistry*, Technical Report. 2018.
2. Alberi K, Nardelli MB, Zakutayev A, Mitas L, Curtarolo S, Jain A, Fornari M, Marzari N, Takeuchi I, Green ML, Kanatzidis M, Toney MF, Butenko S, Meredig B, Lany S, Kattner U, Davydov A, Toberer ES, Stevanovic V, Walsh A, Park N-G, Aspuru-Guzik A, Tabor DP, Nelson J, Murphy J, Setlur A, Gregoire J, Li H, Xiao R, Ludwig A, Martin LW, Rappe AM, Wei S-H, Perkins J: *The 2019 materials by design roadmap*. *J Phys D: Appl Phys* 2018, **52**:013001.
3. Rupp M, Von Lilienfeld OA, Burke K: **Guest editorial: special topic on data-enabled theoretical chemistry**. *J Chem Phys* 2018, **148**:241401.
4. Weininger D: **SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules**. *J Chem Inf Model* 1988, **28**:31-36.
5. Bartók AP, Kondor R, Csányi G: **On representing chemical environments**. *Phys Rev B* 2013, **87**:184115.
6. Kearnes S, McCloskey K, Berndl M, Pande V, Riley P: **Molecular graph convolutions: moving beyond fingerprints**. *J Comput Aided Mol Des* 2016, **30**:595-608.
7. Hy TS, Trivedi S, Pan H, Anderson BM, Kondor R: **Predicting molecular properties with covariant compositional networks**. *J Chem Phys* 2018, **148**:241745.
8. Schütt KT, Arbabzadah F, Chmiela S, Müller KR, Tkatchenko A: **Quantum-chemical insights from deep tensor neural networks**. *Nat Commun* 2016;6:13.
9. Hachmann J, Olivares-Amaya R, Atahan-Evrenk S, Amador-Bedolla C, Sánchez-Carrera RS, Gold-Parker A, Vogt L, Brockway AM, Aspuru-Guzik A: **The Harvard Clean Energy Project: large-scale computational screening and design of organic photovoltaics on the world community grid**. *J Phys Chem Lett* 2011, **2**:2241-2251.
10. Hachmann J, Olivares-Amaya R, Jinich A, Appleton AL, Blood-Forsythe MA, Seress LR, Román-Salgado C, Trepte K, Atahan-Evrenk S, Er S, Shrestha S, Mondal R, Sokolov A, Bao Z, Aspuru-Guzik A: **Lead candidates for high-performance organic photovoltaics from high-throughput quantum chemistry – the Harvard Clean Energy Project**. *Energy Environ Sci* 2014, **7**:698-704.
11. Gong Z, Wu Y, Wu L, Sun H: **Predicting thermodynamic properties of alkanes by high-throughput force field simulation and machine learning**. *J Chem Inf Model* 2018, **58**:2502-2516.
12. Pyzer-Knapp EO, Suh C, Gómez-Bombarelli R, Aguilera-Íparragirre J, Aspuru-Guzik A: **What is high-throughput virtual screening? A perspective from organic materials discovery**. *Annu Rev Mater Res* 2015, **45**:195-216.
13. Butler KT, Davies DW, Cartwright H, Isayev O, Walsh A: **Machine learning for molecular and materials science**. *Nature* 2018, **559**:547-555.
14. Li Z, Omidvar N, Chin WS, Robb E, Morris A, Achenie L, Xin H: **Machine-learning energy gaps of porphyrins with molecular graph representations**. *J Phys Chem A* 2018, **122**:4571-4578.
15. Ramakrishnan R, Dral PO, Rupp M, Von Lilienfeld OA: **Quantum chemistry structures and properties of 134 kilo molecules**. *Sci Data* 2014, **1**:1-7.
16. Ferré G, Haut T, Barros K: **Learning molecular energies using localized graph kernels**. *J Chem Phys* 2017, **146**:114107.
17. Collins CR, Gordon GJ, Von Lilienfeld OA, Yaron DJ: **Constant size descriptors for accurate machine learning models of molecular properties**. *J Chem Phys* 2018, **148**:241718.
18. LeCun Y, Cortes C, Burges CJC: *The MNIST Database of Handwritten Digits*. 2013 <http://yann.lecun.com/exdb/mnist/>.
19. Faber FA, Hutchison L, Huang B, Gilmer J, Schoenholz SS, Dahl GE, Vinyals O, Kearnes S, Riley PF, Von Lilienfeld OA: **Prediction errors of molecular machine learning models lower than hybrid DFT error**. *J Chem Theory Comput* 2017, **13**:5255-5264.
20. Sanchez-Lengeling B, Aspuru-Guzik A: **Inverse molecular design • using machine learning: generative models for matter engineering**. *Science* 2018, **361**:360-365.

This review discusses the recent advances in inverse molecular design using deep generative models.

21. Sanchez-Lengeling B, Outeiral C, Guimaraes GL, Aspuru-Guzik A: **Optimizing Distributions Over Molecular Space. An Objective-Reinforced Generative Adversarial Network for Inverse-Design Chemistry (ORGANIC)**. 2017. Prepr., https://chemrxiv.org/articles/ORGANIC_1_pdf/5309668.

22. Putin E, Asadulaev A, Ivanenkov Y, Aladinskiy V, Sanchez-Lengeling B, Aspuru-Guzik A, Zhavoronkov A: **Reinforced adversarial neural computer for de novo molecular design.** *J Chem Inf Model* 2018, **58**:1194-1204.

23. Kadurin A, Aliper A, Kazennov A, Mamoshina P, Vanhaelen Q, Khrabrov K, Zhavoronkov A: **The cornucopia of meaningful leads: applying deep adversarial autoencoders for new molecule development in oncology.** *Oncotarget* 2016, **8**:10883-10890.

24. Popova M, Isayev O, Tropsha A: **Deep reinforcement learning for de novo drug design.** *Sci Adv* 2018, **4**.

25. Gómez-Bombarelli R, Wei JN, Duvenaud D, Hernández-Lobato JM, Sánchez-Lengeling B, Sheberla D, Aguilera-Iparragirre J, Hirzel TD, Adams RP, Aspuru-Guzik A: **Automatic chemical design using a data-driven continuous representation of molecules.** *ACS Cent Sci* 2018, **4**:268-276.

26. Segler MHS, Kogej T, Tyrchan C, Waller MP: **Generating focused molecule libraries for drug discovery with recurrent neural networks.** *ACS Cent Sci* 2018, **4**:120-131.

27. Goh GB, Siegel C, Vishnu A, Hodas NO: **Using rule-based labels for weak supervised learning: a ChemNet for transferable chemical property.** In: *Prediction Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* London, UK: ACM; 2018, 302-310 <http://dx.doi.org/10.1145/3219819.3219838>.

28. Sultan MM, Pande VS: **Transfer learning from Markov models leads to efficient sampling of related systems.** *J Phys Chem B* 2018, **122**:5291-5299.

29. Afzal MAF, Cheng C, Hachmann J: **Combining first-principles and data modeling for the accurate prediction of the refractive index of organic polymers.** *J Chem Phys* 2018, **148**:241712.

30. Ribeiro JML, Bravo P, Wang Y, Tiwary P: **Reweighted autoencoded variational Bayes for enhanced sampling (RAVE).** *J Chem Phys* 2018, **149**:072301.

31. Wei C, Tan AR, Ferguson AL: **Collective variable discovery and enhanced sampling using autoencoders: innovations in network architecture and error function design.** *J Chem Phys* 2018, **149**:072312.

32. Khorshidi A, Peterson AA: **Amp: a modular approach to machine learning in atomistic simulations.** *Comput Phys Commun* 2016, **207**:310-324.

33. Botu V, Batra R, Chapman J, Ramprasad R: **Machine learning force fields: construction, validation, and outlook.** *J Phys Chem C* 2017, **121**:511-522.

34. Smith JS, Isayev O, Roitberg AE: **ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost.** *Chem Sci* 2017, **8**:3192-3203.

In this study, deep learning is used to fit interatomic potentials and develop the so-called ANI model for transferable data-derived potentials with comparable accuracy to the reference DFT calculations.

35. Chmiela S, Tkatchenko A, Sauceda HE, Poltavsky I, Schütt KT, Müller K-R: **Machine learning of accurate energy-conserving molecular force fields.** *Sci Adv* 2017, **3**.

36. Brockherde F, Vogt L, Li L, Tuckerman ME, Burke K, Müller K-R: **Bypassing the Kohn-Sham equations with machine learning.** *Nat Commun* 2017, **8**:872.

37. Coley CW, Green WH, Jensen KF: **Machine learning in computer-aided synthesis planning.** *Acc Chem Res* 2018, **51**:1281-1289.

This study surveys the role of ML in synthesis planning and the prediction of reaction outcomes.

38. Lowe DM: **Patent Reaction Extraction: Downloads.** 2014 In: <https://bitbucket.org/dan2097/patent-reaction-extraction/downloads>.

39. Nam J, Kim J: **Linking the Neural Machine Translation and the Prediction of Organic Chemistry Reactions.** 2016. Prepr., <https://arxiv.org/abs/1612.09529>.

40. Schwaller P, Gaudin T, Lányi D, Bekas C, Laino T: **“Found in Translation”: predicting outcomes of complex organic chemistry reactions using neural sequence-to-sequence models.** *Chem Sci* 2018, **9**:6091-6098.

41. Coley CW, Jin W, Rogers L, Jamison TF, Jaakkola TS, Green WH, Barzilay R, Jensen KF: **A graph-convolutional neural network model for the prediction of chemical reactivity.** *Chem Sci* 2019, **10**:370-377.

42. **IBM RXN for Chemistry.** 2018 In: <https://rxn.res.ibm.com>.

43. **ASKCOS.** 2018 In: <http://askcos.mit.edu>.

44. Klucznik T, Mikulák-Kluczník B, McCormack MP, Lima H, Szymkuć S, Bhowmick M, Molga K, Zhou Y, Rickershauser L, Gajewska EP, Touthkine A, Dittwald P, Startek MP, Kirkovits GJ, Roszak R, Adamski A, Sieredzińska B, Mrksich M, Trice SLJ, Grzybowski BA: **Efficient syntheses of diverse, medicinally relevant targets planned by computer and executed in the laboratory.** *Chem* 2018, **4**:522-532.

45. Häse F, Roch LM, Kreisbeck C, Aspuru-Guzik A: **PHOENICS: a universal deep bayesian optimizer.** *ACS Cent Sci* 2018, **4**:1134-1145.

46. Tran K, Ulissi ZW: **Active learning across intermetallics to guide discovery of electrocatalysts for CO₂ reduction and H₂ evolution.** *Nat Catal* 2018, **1**:696-703.

47. Gubaev K, Podryabinkin EV, Hart GLW, Shapeev AV: **Accelerating high-throughput searches for new alloys with active learning of interatomic potentials.** *Comput Mater Sci* 2019, **156**:148-156.

48. Hachmann J, Afzal MAF, Haghigatlari M, Pal Y: **Building and deploying a cyberinfrastructure for the data-driven design of chemical systems and the exploration of chemical space.** *Mol Simul* 2018, **44**:921-929.

This paper presents a software ecosystem for the development and broader dissemination of techniques at the different stages of a molecular data mining workflow.

49. Krylov A, Windus TL, Barnes T, Marin-Rimoldi E, Nash JA, Pritchard B, Smith DGA, Altarawy D, Saxe P, Clementi C, Crawford TD, Harrison RJ, Jha S, Pande VS, Head-Gordon T: **Perspective: computational chemistry software and its advancement as illustrated through three grand challenge cases for molecular science.** *J Chem Phys* 2018, **149**:180901.

50. Wilkins-Diehr N, Crawford TD: **NSF's inaugural software institutes: the science gateways community institute and the molecular sciences software institute.** *Comput Sci Eng* 2018, **20**:26-38.