



Published in final edited form as:

*Expert Opin Drug Discov.* 2021 September ; 16(9): 1009–1023. doi:10.1080/17460441.2021.1925247.

## Assigning Confidence to Molecular Property Prediction

AkshatKumar Nigam<sup>1,2</sup>, Robert Pollice<sup>1,2</sup>, Matthew F. D. Hurley<sup>1,3</sup>, Riley J. Hickman<sup>1,2</sup>, Matteo Aldeghi<sup>1,2,4</sup>, Naruki Yoshikawa<sup>1,2</sup>, Seyone Chithrananda<sup>1</sup>, Vincent A. Voelz<sup>1,3</sup>, Alán Aspuru-Guzik<sup>1,2,5,6,\*</sup>

<sup>1</sup>Department of Computer Science, University of Toronto, Canada.

<sup>2</sup>Department of Chemistry, University of Toronto, Canada.

<sup>3</sup>Department of Chemistry, Temple University, Philadelphia, PA 19122, USA.

<sup>4</sup>Vector Institute for Artificial Intelligence, 661 University Ave Suite 710, Toronto, Ontario M5G 1M1, Canada.

<sup>5</sup>Vector Institute for Artificial Intelligence, Toronto, Canada.

<sup>6</sup>Lebovic Fellow, Canadian Institute for Advanced Research (CIFAR), 661 University Ave, Toronto, Ontario M5G, Canada.

### Abstract

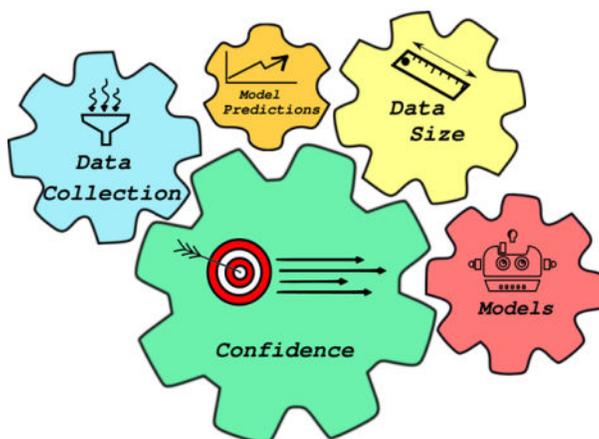
**Introduction:** Computational modeling has rapidly advanced over the last decades, especially to predict molecular properties for chemistry, material science and drug design. Recently, machine learning techniques have emerged as a powerful and cost-effective strategy to learn from existing datasets and perform predictions on unseen molecules. Accordingly, the explosive rise of data-driven techniques raises an important question: What confidence can be assigned to molecular property predictions and what techniques can be used for that purpose?

**Areas covered:** In this work, we discuss popular strategies for predicting molecular properties relevant to drug design, their corresponding uncertainty sources and methods to quantify uncertainty and confidence. First, our considerations for assessing confidence begin with dataset bias and size, data-driven property prediction and feature design. Next, we discuss property simulation via molecular docking, and free-energy simulations of binding affinity in detail. Lastly, we investigate how these uncertainties propagate to generative models, as they are usually coupled with property predictors.

**Expert opinion:** Computational techniques are paramount to reduce the prohibitive cost and timing of brute-force experimentation when exploring the enormous chemical space. We believe that assessing uncertainty in property prediction models is essential whenever closed-loop drug design campaigns relying on high-throughput virtual screening are deployed. Accordingly, considering sources of uncertainty leads to better-informed experimental validations, more reliable predictions and to more realistic expectations of the entire workflow. Overall, this increases confidence in the predictions and designs and, ultimately, accelerates drug design.

### Graphical Abstract

\*Correspondence to: alan@aspuru.com.



## Keywords

Neural networks; deep learning; drug discovery; generative models; artificial intelligence; model uncertainty estimation; docking; molecular dynamics; generative models

---

## I. INTRODUCTION

While data-driven modeling has made large advances in image processing and speech recognition, ground-breaking contributions in drug discovery are harder to find. Often, only limited amounts of reliable *in vitro* and *in vivo* data suitable for supervised learning tasks are available. Additionally, uncertainties of experimental data, for instance regarding effectiveness and toxicity, are common and significant as the experimental design is complex, prone to noise and the outcome is dependent on many parameters, in particular compound dosage. Hence, these uncertainties are propagated to property prediction workflows when data-driven models are trained on that data.

While it is evident that machine learning (ML) prediction accuracy and data quality are intrinsically related, there are more subtle aspects that are sometimes neglected. Dataset size, composition, coverage and training-test split have considerable influence on the final model performance but are insufficient to assess uncertainty alone without suitable models and methodologies for that purpose. For instance, using random or scaffold-based train/test splits does not lead to reliable measures of model performance and generalizability. Other common problems include high dataset bias, small sample size or low chemical diversity. However, some of these issues currently do not have a straightforward remedy as there is no commonly accepted way to describe chemical subspaces rendering it difficult to uncover all inherent biases. Nevertheless, when the underlying data is not investigated appropriately model performance and generalizability tend to be overestimated leading to a significant number of false predictions and, ultimately, low confidence in the property prediction workflows.

In this article, we discuss methodologies and datasets for chemical properties important in drug design. In particular, we focus on the main sources of uncertainties for both simulation-based and data-driven property prediction. In doing so, we discuss uncertainties inherent in

datasets, outputs of data-driven models, input features, and simulation of binding affinities. In addition, we discuss the importance of uncertainty in generative models, especially for property-based molecule design. Finally, we close with our personal opinion on the most important aspects of uncertainty and confidence focusing on important problems and future avenues that will lead to higher predictive ability and models that naturally take uncertainty into account.

## II. DATASET UNCERTAINTY

One of the most common challenges for data-driven molecular property prediction is dataset size and composition. For many important properties in drug discovery, especially pharmacologic properties of molecules such as absorption, distribution, metabolism, excretion and toxicity (ADMET), only a limited amount of high-quality data is available and it is usually only available for certain classes of molecules, which inherently introduces biases. To provide an overview of the amount and type of data available, Table 1 lists popular datasets for molecular property prediction. Accordingly, in this section, we discuss dataset characteristics that are important to consider for data-driven property prediction tasks along with methods and procedures to minimize their impact on property prediction performance.

In the framework of supervised learning, models learn relationships between input and output from the training data and, subsequently, predict properties for unlabelled input data, and the corresponding prediction ability is termed generalization. When predictors perform well on the training set, but performance on the unlabelled data is poor, the model is overfitted. This is a very common problem in ML, especially when the training set size is small (Figure 1-A). Learning theory states that collecting more data improves generalization if the data is independent and identically distributed [20]. However, real data is not a uniform sample of chemical space. Typically, molecules in a dataset are collected under specific criteria such as the number of atoms, the constituent elements, the similarity to known molecules or the availability of synthetic procedures introducing bias (Figure 1-B). In addition, there is an inherent bias in both industry and academia to publish only successful experiments which can be problematic for assembling training data as negative results are as important as positive ones for data-driven modeling.

First, the uncertainty of model predictions for unlabelled molecules is largely dependent on the difference of the structure and property distributions to the training data. To quantify this difference, in the field of quantitative structure-activity relationships (QSARs) [21], the concept of applicability domain (AD) is widely applied. The AD of a QSAR model is defined as “the response and chemical structure space in which the model makes predictions with a given reliability”[22] and makes the inherent bias of a model comprehensible (Figure 1-C). There are several methods to estimate ADs [23]. A common approach is to consider molecules with descriptors within a certain distance to the mean of the training data to be inside the AD. Notably, predictions for molecules inside the AD are considered reliable, outside the AD reliability cannot be guaranteed.

Secondly, a significant bias in the training data can cause models only to learn the inherent bias rather than physically meaningful relationships (Figure 1-B). When users are unaware of this bias, it inadvertently leads to overly optimistic conclusions concerning model performance. Recently [10], it was pointed out that there is significant hidden bias in the widely used dataset for structure-based virtual screening, the Directory of Useful Decoys: Enhanced (DUD-E)[24]. DUD-E includes two types of molecules: ligands, small molecules that bind to a receptor, and decoys, molecules with similar physical properties but dissimilar chemical structures compared to the ligands. Using DUD-E as a dataset for training, the performance of receptor-ligand and ligand-only ML models were compared to determine whether the receptor-ligand models also learned from the actual interactions or whether they just learned the inherent ligand bias. Ideally, the performance of these two types of models should be different as the binding energies are a function of both ligand and receptor structures, and the receptor structure can generally not be inferred from the ligand alone. However, the performance was found to be equivalent, the receptor-ligand models did not learn from the receptor structure, and it was concluded that "models may only learn the inherent bias in the dataset rather than physically meaningful features" [10]. This calls for general methods to test what property predictors actually learn as it would naturally lead to more powerful, generalizable models and to more confidence in the models. To minimize hidden biases, the importance of dataset diversity for the generalizability of ML models has been pointed out before [25].

Finally, small datasets pose a particular challenge as they are inherently biased due to their limited size and, additionally, overfitting is to be expected. An important strategy to reduce both bias and overfitting in ML models when data is sparse is transfer learning (Figure 1-D) [23, 26]. When using that approach, models are first pre-trained on large, sometimes unlabelled, datasets, and subsequently trained on the property of interest. Importantly, transfer learning has been shown to improve the prediction performance for models trained on small datasets significantly [27, 28]. For instance, recently [27], a language-based ML property prediction model was pre-trained first via self-supervised learning on one million unlabeled molecules curated from ChEMBL. This pre-training improved the prediction performance after subsequent training on small datasets of molecular activity consistently, regardless of the specific property prediction task and the training dataset size. Notably, the smallest dataset only contained 642 samples [27].

### III. UNCERTAINTY IN THE OUTPUTS

Supervised ML largely relies on an inductive inference approach to derive general rules from a finite number of training examples. Therefore, predictive models obtained using inductive inference are never formally correct and inherently uncertain. There are often additional sources of uncertainty that stem from noise or imprecision on target measurements. For instance, molecular properties like binding affinity are often measured with an uncertainty, e.g.  $1 \mu M \pm 10\%$ . The latter type is referred to as aleatoric uncertainty, while the former is referred to as epistemic or systematic uncertainty and can in theory be reduced in light of more training observations or increased knowledge about the correct model. Both types of uncertainty are important, although they are often not treated separately and the formal distinction between the two remains an active area of research [29]. Many ML methods

are available for representing uncertainty. They can be categorized with respect to the way uncertainty is represented and whether or not they allow for differentiation between the two types (aleatoric and epistemic). Here we divide the discussion of methods derived from frequentist statistics and ones based on Bayesian inference.

A widely used class of frequentist methods are ensemble strategies which typically use a large set of distinct predictors as opposed to a single model. Ensembles of tree-based models such as random forests [30] are commonly used to produce frequentist probability estimates based on the relative output frequency of the ensemble members [31, 32]. Ensemble learning can introduce variance through multiple ensemble models with different parameter initializations (ensembling) or through the use of randomly sampled training sets (bagging or boosting, Figure 2-a). Ensembles of neural networks can also provide computationally efficient, readily parallelizable uncertainty estimates [33, 34]. However, care must be taken using these methods as the uncertainties are often biased, poorly calibrated, and are only concerned with aleatoric uncertainty. Heteroscedastic aleatoric uncertainty can also be modeled using a single neural network trained with frequentist maximum likelihood inference that outputs a probability distribution (Figure 2-b) [35]. Conformal prediction (CP) has recently received significant attention, in which one ML model is trained to predict a property and another to predict the uncertainty (Figure 2 d) [36–39]. CP methods are based on a rigorous mathematical framework and allow users to select intuitive confidence levels for their predictions, e.g. selecting a confidence level of 0.9 means that at most 10% of predictions will be outside the predicted range.

Methods inspired by Bayesian inference seek to update a prior distribution  $p(\theta)$  on the space of all possible models in light of training data,  $D$ , yielding a posterior distribution  $p(\theta|D)$ . Given a molecular representation for prediction,  $x^*$ , the predictive posterior distribution  $p(y^*|x^*, D)$  is in theory obtained by averaging all possible models weighted by their respective posteriors. Bayesian methods have experienced a recent resurgence in drug design applications due to increasing computational power and advancement in algorithms for approximate inference.

Gaussian processes (GPs) [40] allow for Bayesian inference in a non-parametric way, where the prior is determined by a mean and kernel function. The prior may be updated with observed data and for regression with Gaussian noise, the predictive posterior distribution is Gaussian with mean  $\mu$  and variance  $\sigma^2$ . The predictive variance  $\sigma^2$  represents the total uncertainty but can be decomposed into the variance of the error term  $\sigma_e^2$  which represents aleatoric uncertainty. Epistemic uncertainty is the difference between  $\sigma^2 - \sigma_e^2$  and is determined by the hyperparameters of the kernel function such as the characteristic length scale. GPs are commonly used tools for property prediction in the low data regime (i.e. below 1000 data points) due to their known robustness to overfitting [40–42]. Hie et al. [43] used a GP trained on less than 100 compounds to screen a large library for candidates with nanomolar affinity for diverse kinases and whole-cell growth inhibition of *Mycobacterium tuberculosis* [43]. The authors show how uncertainty estimation can enable machine-guided discovery. However, the traditional implementation of GPs incurs cubic scaling with training set size, rendering them less desirable when training data is abundant.

Bayesian neural networks (BNNs) place probability distributions on neural network parameters  $w$  and the analytically intractable posterior distribution  $p(w|D)$  is typically estimated using variational Bayesian inference [44]. The predictive posterior distribution is approximated using Monte Carlo sampling, which intuitively involves averaging the predictions of multiple parameter configurations sampled from the posterior (Figure 2-c). Replacing point-estimated network parameters with distributions allows for the quantification of epistemic uncertainty as well as heteroscedastic aleatoric uncertainty. Recently, such networks have been combined with state-of-the-art representation learning to give uncertain predictions of the physiological properties of small molecules [45].

For classification tasks, confidence calibration seeks to produce probability estimates which correctly reproduce the true correctness likelihood. In other words, given  $N$  predictions each with confidence 0.9 about a label, we would expect 90% of the predictions to give the correct label. Scalar valued confidence statistics such as expected calibration error and maximum calibration error have also been proposed. Many calibration methods such as isotonic regression or temperature scaling are based on a learned post-processing step in which a calibrated probability  $q_i$  is derived from the model's output probability  $p_i^*$  [46]. Confidence calibration schemes can be applied as a scalable post-processing step to any predictive learning algorithm, and their ideas can be extended to regression tasks.

Confidence or uncertainty estimation in molecular property prediction is crucial for tasks that involve data-driven decision making where valuable resources or human risk is at stake. For example, BNNs have been employed to give intuitive probabilistic statements about the severity of the risk of drug-induced liver injury [47, 48]. Predicting clinical outcomes from pre-clinical data and using uncertainty estimates to quantify and communicate risks can preserve time, resources, and patient well-being [49]. Predictive uncertainty is central to the field of active learning [50], which has been known in the context of drug discovery for almost two decades. Active learning strategies balance exploitation of the knowledge of a predictive model to retrieve active candidate compounds with the exploration of regions in compound space where the model is highly speculative [45, 51, 52]. Explorative behaviour of an active learner is paramount when structural novelty is prioritized. Elucidation of error sources in training data remains a critical application for uncertain predictive models. Ryu et al. modeled aleatoric and epistemic uncertainty separately using Bayesian graph convolutional neural networks for a variety of molecular property prediction tasks including classification of bioactivity and toxicity. Analysis of predictive aleatoric uncertainty allowed for the identification of poor quality training data for some molecular candidates [53].

In the near future, we believe predictive uncertainty will become an indispensable component of molecular property prediction. Researchers will continue to systematically benchmark uncertainty quantification methods on various datasets [54]. We anticipate software packages that facilitate uncertain predictions will continue to appear [55].

#### IV. UNCERTAINTY IN THE INPUT FEATURES

Contrary to uncertainty in the properties being predicted, imprecision in the input features is a source of uncertainty that is rarely considered. Property prediction in drug design often

relies on representations encoding a specific molecule, such as SMILES strings [58], or on three-dimensional structural data. However, these static descriptions of molecular systems are not representative of their dynamic nature in solution. In this section, we discuss how input-related uncertainty might arise when predicting chemical properties from molecular structures, and how some of the latest ML models might be used to capture this uncertainty in the confidence of the predictions.

One form of input uncertainty is due to the presence of multiple protonation and tautomeric states in solution. String and graph representations of molecules are static and unable to capture these multiple states and their probabilities in solution. SMILES strings representing two different tautomers are likely to return different property predictions when they in fact should be the same given the behaviour of these molecules in solution. This represents a challenge to many compound property predictions and cheminformatics tools and may be seen as a form of input uncertainty, given that the chemical species in solution can be modeled only approximately. Similarly, protonation states depend on the chemical environment of the molecule, there may not be one dominant state in solution and, finally, these may change upon interaction with a binding partner. All these aspects contribute to creating uncertainty about the input representation used in ML prediction tasks.

Three-dimensional structural data is often used to predict properties like drug binding affinity, selectivity, and mutational effects on protein stability and drug resistance. This data is typically derived by X-ray crystallography and, increasingly, cryo-electron microscopy (cryo-EM). Despite providing abundant information on protein-ligand interactions, which can be used to drive drug design, these structures are only a single snapshot of all possible protein-ligand conformations in solution. As such, this three-dimensional data is a single datapoint of a much broader input probability distribution. Some structure-based modeling approaches, like ensemble docking, try to take this uncertainty into account by sampling multiple conformations with molecular dynamics simulations. In such a way, ensemble docking tries to capture the distribution/uncertainty of protein-ligand structures. By generating an “ensemble” of drug target conformations, the docking results are less likely to be overfitted to a specific protein conformation. However, fast scoring functions for ligand binding affinities are typically trained on databases of static structures, such as the PDBbind database, and do not capture this type of uncertainty in their predictions. In addition to the statistical uncertainty caused by conformational variability, differences between protein structures *in crystallo* and *in vivo* due to packing effects and cryo temperatures may result in systematic biases.

Ideally, one would like uncertainty about multiple possible chemical and conformational states in solution to be captured by the ML model and reflected in the confidence of the property predictions. Currently, only a handful of specialized ML approaches are able to capture uncertainty in the input features. Gaussian process latent variable models, together with other similar kernel-based approaches [59, 60], are examples of such models [61, 62]. Recent work has also started to investigate the effect of molecular ensembles in property predictions [63] with ML models. However, we are unaware of widely adopted supervised ML models that can either take into account known input uncertainty, or that can infer it directly from data. We believe this to be an underexplored research direction with much

potential for practical applications. Such models could become an indispensable component of molecular property prediction, being able to capture or infer feature uncertainty in structural and chemical data.

## V. UNCERTAINTY IN THE BINDING AFFINITIES

A key property of interest in drug discovery is the binding affinity of a ligand to a target receptor. Of the many physics-based models that have been developed and refined to predict binding affinity, here we focus on two widely-used techniques that utilize the three-dimensional structures of ligand and receptor, and a molecular-mechanics energy function to compute ligand affinities: molecular docking and free energy perturbation (FEP) methods.

Of the two, molecular docking is less computationally expensive and more suitable for high-throughput virtual screening of potential inhibitors. Docking algorithms such as AutoDock Vina [64], Glide [65], Surflex-Dock [66] and GOLD [67] can predict correct bound poses ( $< 2 \text{ \AA}$  RMSD) with an accuracy of 70–80% [68] and can correctly rank (trios of low-, medium- and high-affinity) ligands with an accuracy of almost 60% [68–70]. Predictions of absolute binding affinity from docking scores, however, generally show a poor correlation to experimental binding affinities [71], with  $\text{pK}_a$  (decadic logarithm of the affinity) prediction accuracies of 1.5–2.0 (RMSE) [69, 72].

Large uncertainties in docking estimates can be attributed to conceptual deficiencies in the scoring function, for example, a lack of proper accounting of solvation thermodynamics, which may require incorporation of explicit water molecules against an ensemble of receptor poses to better represent the physical system [73]. Docking scores are also made uncertain by implementation choices made for algorithmic efficiency. For example, flexible ligand or receptor degrees of freedom are usually restricted to enable efficient conformational searching. Recent studies using deep learning neural-network-based scoring functions show promise in reducing prediction uncertainties [72, 74]. Regardless of these inaccuracies, docking approaches are likely to remain an essential component of virtual screening, as they can reduce and enrich the chemical space of potential binders, and provide initial bound-pose conformations for seeding molecular dynamics (MD) simulations.

In contrast to docking, FEP methods are much more expensive, requiring all-atom MD simulations to perform statistical sampling of ligand-receptor configurations [75–78]. The benefit is a best-in-class prediction of ligand binding affinity, achieving accuracies within chemical accuracy, i.e. 1.0 kcal/mol [79, 80]. Assigning uncertainties to such predictions, however, requires consideration of the many error sources these methods introduce. These can be categorized as arising from both sampling and scoring inaccuracies.

### A. Uncertainty in FEP Estimates due to Sampling Error

To understand uncertainty from statistical sampling error, consider that using MD simulation naïvely to predict ligand affinity would require equilibrium sampling of multiple reversible binding transitions which typically exceed the timescales of state-of-the-art MD, even on special-purpose high-performance computers [81]. FEP and related methods circumvent this

problem by performing simulations in multiple thermodynamic ensembles for a series of alchemical intermediates and analyzing the resulting samples with a statistical free energy estimator. Absolute binding free energy (ABFE) is computed along alchemical paths that decouple all non-bonded interactions in solution of the ligand, including the binding to a receptor. Relative binding free energy (RBFE) is computed along with an alchemical transformation from one molecular topology to another.

Statistical free energy estimators that utilize multiensemble sampling include exponential averaging (EXP) [82], thermodynamic integration (TI) [83], the weighted-histogram method (WHAM) [84], the Bennett acceptance ratio (BAR) method [85], the multistate Bennett acceptance ratio (MBAR) method [86], and newer dynamical estimators like TRAM [87] and DHAM [88]. Uncertainties from statistical sampling can influence the confidence of calculated Gibbs free energies in complicated ways. Hence, regardless of which free estimator is used, a non-parametric bootstrap is preferred to estimate uncertainty in Gibbs free energies [76]. Similarly, uncertainties can be calculated over block-averages of a long and equilibrated simulation trajectory, or can alternatively be estimated across multiple independent simulations, in some cases initiated with different ligand starting poses. For FEP, which involves computing free energy differences between the ligand-bound state and the free ligand in solution, overall uncertainties are often obtained by summing individual uncertainties for each ensemble in quadrature [71, 76].

Different free energy estimators have unique bias and variance due to finite sampling, but they all depend on having sufficient thermodynamic overlap between neighbouring ensembles along the thermodynamic path. For this reason, relative free energy estimation is often more accurate for pairs of similar molecules, since the alchemical transformation involves a smaller number of atoms, and hence inherently has a better overlap. Achieving good thermodynamic overlap is the main consideration guiding choices about the number and type of alchemical intermediates, which may require position or angular restraints on the ligand to focus sampling on particular ligand poses [71].

Most free-energy estimators assume that samples are statistically independent and drawn from the equilibrium distribution. Sufficient equilibration can be notoriously slow to achieve for some systems, e.g. ligands with “hidden barriers” between conformational states, a problem which Hamiltonian replica exchange and other enhanced sampling methods have been designed to address [16]. To avoid statistically correlated input, which results in artificially low uncertainty estimates, recommended strategies discard initial unequilibrated samples, and subsample snapshots at a frequency  $(2\tau + 1)$ , where  $\tau$  is a computed correlation time [89]. An algorithm to automatically detect simulation equilibration in this way has been developed [90].

Related to FEP is a class of non-equilibrium work (NEW) free energy methods, based on the Crooks fluctuation theorem [91], which drives a system between alchemical endpoints and analyzes the distribution of non-equilibrium work values [77]. These methods can achieve comparable accuracies to FEP with fewer intermediates [92–94], yielding uncertainties that depend on the variance of the work distributions, in turn depending on the timescale of the alchemical schedule. Sampling uncertainty can also be introduced by artifacts from

the implementation of MD. For example, uncertainties on thermodynamic and kinetic properties can depend greatly on solvation artifacts that depend on simulation box size [95]. Furthermore, surprisingly, variation in computed energies across different simulation packages [96], integrators, and thermostats may be non-negligible [97].

## B. Uncertainty in FEP Estimates due to Scoring Error

The workhorse of FEP is a molecular mechanics energy function, called a force field, which has usually been designed to approximate a quantum mechanical potential energy surface, and tuned to reproduce experimental observables. Iterative cycles of force field development, parameterization and assessment are thus vital for reducing the uncertainty introduced by force field inaccuracy [94, 98, 99]. With new machine learning advances have come new ways to improve force fields for the prediction of ligand binding free energies. One product of the recently-launched Open Force Field Initiative is a new generation of force fields that eschews atom-types in favour of direct chemical perception, utilizing the SMIRKS-native Open Force Field (SMIRNOFF) scheme for parameter assignment [100]. Parameterization of Open Force Field v1.0.0, code-named Parsley [38], was made possible by the ForceBalance algorithm, a Bayesian algorithm for automatic and reproducible parameterization [101–103]. Polarizable force fields such as AMOEBA continue to make gains in accuracy for predictions of Gibbs free energy of binding [104, 105], and there are examples of neural network-based potentials, which aim to reproduce molecular properties with DFT accuracy at force field cost [106, 107].

Finally, uncertainty can also result from the assumptions of a selected model. For example, most fixed-charge force fields require the pre-selection of particular protonation states for both ligand and ionizable side chains of a receptor. Depending on the predicted  $pK_a$  of acidic or basic groups, which may differ significantly between bound and unbound states, multiple simulations may be required to explore these possibilities. Overall, there are several factors to consider when designing a study to compute binding affinities for drug-like molecules. New methodologies and best practices can often be found through blind assessment challenges such as SAMPL [108] (Statistical Assessment of Modeling of Proteins and Ligands), which are created annually to assess predictions of binding affinities among other properties important to drug discovery.

## VI. UNCERTAINTY IN GENERATIVE MODELING

Generative models can produce synthetic data samples, some of them by learning from reference data. For molecular design, generative models are used to tackle the inverse design problem [109], where novel structures with desired properties are generated. Molecular property estimators are often used in conjunction with these models [110]. In this section, we discuss the importance of property prediction uncertainty in the context of generative modeling. In particular, we consider variational autoencoders (VAEs) [111], generative adversarial networks (GANs) [112], reinforcement learning (RL) [113], and genetic algorithms (GAs) [114].

VAEs possess the ability to take discrete structures such as molecules and convert them into high-dimensional latent representations by way of neural networks (referred to as encoders).

Simultaneously, a mapping from the latent representation back to the discrete structure is trained using a separate neural network (referred to as a decoder). During training, VAEs learn to encode and decode correct molecules and simultaneously shape the latent space into a gaussian distribution for sampling. For molecular design, Gómez-Bombarelli et al. [115] introduced an additional neural network that predicts molecular properties from the latent space representation. During the training of this network, gradient optimization steps are taken for both the parameters of the encoder network and the property prediction network. This shapes the latent space based on the predicted property.

The formulation of VAEs as generative models brings about three sources of uncertainty for the associated property prediction model: (1) Decoding points from a latent space does not always lead to valid molecules. However, for every point in the latent space one can predict molecular properties via the prediction model. Hence, these regions can be viewed as uncertain spaces where predictions have to be disregarded due to the lack of a physical manifestation. (2) The encoding of structures to the latent space is stochastic. Kingma et al. [111] introduced a reparameterization step that ensures differentiability through the latent space for training the overall architecture and imposition of a gaussian distribution. As a result, identical molecules are typically encoded at a range of points in the latent space. While these encoded molecules should have the exact same predicted property, the range of positions tends to give rise to different predictions. Hence, there are multiple property predictions associated with any given molecule representing the corresponding uncertainty. (3) Discrete objects within VAEs, in particular molecular string representations and adjacency matrices, are usually decoded from the latent space as a time-series. These discrete objects are generated one at a time, by sampling over the distribution of all user-defined actions, i.e. units such as the allowed characters in a string. Typically, the unit with the largest probability is selected. However, multiple units can have identical probabilities. Consequently, decoding the same point from the latent space multiple times can result in distinct molecules. Additionally, as the same latent space point represents all these different molecules simultaneously, they will have the exact same predicted property.

GANs [112, 116, 117] are generative models with joint training of two competing networks, a generator and a discriminator. Like VAEs, they rely on a gaussian latent space to draw samples from and decode them using a neural network (referred to as generator network). The generated samples are compared to references from a dataset by another neural network (referred to as discriminator). The generator and discriminator compete with each other, the generator tries to produce samples that appear more realistic to the discriminator while the discriminator tries to improve its reliability to distinguish generated from reference samples. Similar to VAEs, in the traditional formulation of GANs, property predictors can be associated with the latent space. The sources of uncertainty are essentially identical to VAEs.

RL is a field where agents learn to navigate environments to maximize rewards. In molecular design, RL agents build chemical compounds [118–120], and rewards are provided for achieving desired molecular properties. Actions constitute making changes to chemical structures, while rewards are calculated on-the-fly based on the resulting structures. In contrast, GAs [121–123] are generative models that generate structures without gradient-based optimization. The algorithm is initiated with a set of molecules, i.e. the initial

population, that are ranked by calculating their fitness. For molecular design, the fitness is typically assessed based on a set of target properties. The fittest molecules are carried over to the next generation, the unfit molecules are replaced with mutations and crossovers of the best. For both RL and GAs, to offset the often prohibitive computational cost of obtaining accurate property predictions, cheap approximations such as ML models are utilized to estimate rewards and fitness, respectively. By using cheap property predictors, in high-dimensional spaces, the generation of structures can be skewed to uncertain regions of the chemical space. For instance, certain structures can be adversarial examples or outliers for property predictors, where the model can be overly confident or uncertain about its prediction. Hence, the resulting molecule modifications are uncertain due to the uncertainty in the rewards. Notably, Thiede et al. [124] demonstrated that the uncertainty from prediction networks can be used directly in the reward function to guide the generation of structures outside of the confidence region of the prediction network for curiosity-driven molecular design. Interestingly, Shen et al. [125] propose PASITHEA, where a property prediction neural network is directly used for generating molecules with deep dreaming. An interesting direction would be the incorporation of uncertainty within property predictors for structure generation via the deep dreaming algorithm.

Generative models coupled with property predictors are yet to achieve wide adoption due to their complex setup, their high computational costs and the need to adapt them to specific problems. When used for properties that are inherently complicated and hard to predict, the role of uncertainty and prediction confidence will be very important. In addition, in models with latent spaces such as VAEs and GANs, the choice of molecular representation influences both the validity of the latent space and thereby the prediction confidence. For instance, the use of VAEs and GANs was demonstrated in conjunction with the SMILES [58] representation, where the probability to produce invalid strings is high. Consequently, both VAEs and GANs trained on SMILES can produce invalid molecules. Importantly, this problem has recently been solved by the SELFIES representation [126], which guarantees molecular validity by design. While SELFIES increases model confidence by generating only valid molecules, it does not remove uncertainty from inherent model stochasticities such as probabilistic encoding and decoding. Furthermore, confidence in molecular property predictors can be increased by extending datasets used for training and validation of property predictors systematically. Recently, STONED [127], a simple generative model producing molecules at high speed via systematic navigation within SELFIES, has been introduced. Accordingly, based on an initial dataset, it can be used to diversify the molecule space systematically leading to better property prediction models.

## VII. CONCLUSION

In this review, we inspected common sources of uncertainties encountered in data-driven drug design workflows, some of which are well-studied in the literature, and some of which are typically neglected. Hidden biases in data can lead to overly confident property predictions that fail to manifest in the real world. Data-driven prediction models that explicitly incorporate uncertainty convey insights into the confidence of the model itself and provide a handle to assess the prediction quality. Accounting for uncertainty in the input features is underappreciated in ML research but is inherent to complex systems such as

the chemistry of the human body. In contrast, error estimation in the prediction of binding free energies is one of the central methods in the field to assess the quality of simulation workflows that guides the development of efficient sampling techniques. Uncertainties in generative models are another underappreciated topic that needs to be addressed to improve the success of property-driven molecular design approaches. Overall, it is evident that assigning confidence in molecular property predictions is intimately tied to understanding the associated sources of error, and we believe that future progress in this field will require the development and application of even more methods accounting for uncertainties.

## VIII. EXPERT OPINION

Assigning confidence to predictions is directly connected to uncertainty identification and quantification. When training ML models for predicting complex molecular properties, users need to be aware of potential biases of the training data when assessing generalizability. The recent use of DUD-E to train data-driven models serves as a cautionary tale for models learning to exploit biases in the data rather than learning the underlying physics of intermolecular interactions [24]. Most importantly, the smaller the training dataset size the larger the danger of hidden bias and the higher the importance of accounting for them. In that regard, transfer learning, especially in extremely sparse data regimes, requires more research to establish it as the method of choice for molecular property prediction.

Moreover, not only do small datasets lead to significant biases but they also lead to higher epistemic uncertainty because the model does not have enough information to learn the true relationship between input features and output labels [29]. However, expanding datasets, especially for properties relevant in drug design, can be prohibitively expensive and very time-consuming, in particular in the case of experimental data. Furthermore, not every new datum added to an existing dataset provides the same amount of additional information. Accordingly, developing systematic methods to identify data to be added to the existing dataset to maximize the improvement in prediction accuracy and the reduction in epistemic uncertainty is the subject of ongoing research.

Additionally, the aleatoric uncertainty [29] that is inherent in the data used to train ML models is at least as important as the epistemic one. One prominent example from the field of drug discovery discussed in this review is the use of estimated and simulated binding affinities to assess the interaction strength between ligands and receptors. Notably, binding affinities estimated via molecular docking are only of limited accuracy, but their computation efficiency allows assembling large datasets. The main question is whether the estimated property, regardless of how efficient it can be computed, conveys any information about the real relationship between input features and model output. When noise dominates over valuable information, the largest dataset will not provide useful information for ML models to be learned. In summary, it is important to build datasets and construct models that explicitly and consistently account for both aleatoric and epistemic uncertainties. While uncertainty estimates are relatively straightforward to obtain in general, they are only rarely benchmarked or evaluated during model training and future work needs to establish robust procedures for that purpose.

Another important and often underappreciated source of uncertainty, especially in complex systems, are the input features. This is particularly true for chemistry where molecular properties are generally conformer-dependent, and most molecules adopt much more than just one conformation under ambient conditions. Data-driven models, in contrast, are largely trained on 2D molecular representations that neglect conformations entirely. Accordingly, incorporating 3D structures and appropriate features in ML models is an active field of research [63]. The main challenge is to offset the added cost of generating representative 3D conformations with higher prediction accuracy in the final models. Currently, 2D molecular representations are still the state of the art in the field but we believe 3D features will become the standard in the near future.

One topic of paramount importance in drug design is simulating binding affinity between ligands and receptors. For that purpose, docking is a popular technique that provides crude estimates of interaction geometries and strengths using extremely efficient scoring approaches. However, this efficiency comes at the price of accuracy and, hence, cannot match the quality of FEP simulations, which are several orders of magnitude more computationally demanding. Importantly, even among FEP approaches there is an inevitable trade-off between cost and accuracy as the extent of sampling conducted directly controls the quality of the results. Accordingly, for extremely accurate results [71], FEP methods need to capture multiple binding transitions and this can increase the simulation cost by orders of magnitude making it less practical. Finally, a large portion of the uncertainties from both docking approaches and MD simulations stem from the scoring functions and the force fields used, respectively. While systematic optimization procedures for force field parameters are being developed and implemented, the development of force fields that not only provide energy but also uncertainty estimates has received less attention and we believe that future work in that direction is necessary to deliver more accurate models.

Finally, the generative models for molecular design introduced recently rely on efficient molecular property predictors to generate molecules with both desirable structure and function [110]. However, an explicit account and treatment of uncertainties, unfortunately, is largely neglected in the field. While accounting for property prediction uncertainties seems relatively straightforward with current methodologies, dealing with uncertainties in the molecular structure itself, and modifications thereof remains unexplored territory in comparison. Accordingly, we believe that future research in that direction is paramount and has the potential to lead to significant advances in computer-driven molecular design in general and drug design in particular.

## ACKNOWLEDGEMENTS

R.P. acknowledges funding through a Postdoc.Mobility fellowship by the Swiss National Science Foundation (SNSF, Project No. 191127). R.J.H. gratefully acknowledges NSERC for provision of the Postgraduate Scholarships-Doctoral Program (PGSD3-534584-2019). M.A. is supported by a Postdoctoral Fellowship of the Vector Institute. V.A.V and M.F.D.H acknowledge support in part by National Institutes of Health grant 1R01GM123296. A. A.-G. thanks Anders G. Frøseth for his generous support. A. A.-G. also acknowledges the generous support of Natural Resources Canada and the Canada 150 Research Chairs program. We also acknowledge the Department of Navy award (N00014-19-1-2134) issued by the Office of Naval Research. The United States Government has a royalty-free license throughout the world in all copyrightable material contained herein. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Office of Naval Research.

## References

- [1]. Irwin John J, Tang Khanh G, Young Jennifer, Dandarchuluun Chinzorig, Wong Benjamin R, Khurelbaatar Munkhzul, Moroz Yurii S, Mayfield John, and Sayle Roger A. Zinc20—a free ultralarge-scale chemical database for ligand discovery. *Journal of Chemical Information and Modeling*, 2020.
- [2]. Ramakrishnan Raghunathan, Dral Pavlo O, Rupp Matthias, and Von Lilienfeld O Anatole. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1):1–7, 2014.
- [3]. Nakata Maho and Shimazaki Tomomi. Pubchemqc project: a large-scale first-principles electronic structure database for data-driven chemistry. *Journal of chemical information and modeling*, 57(6):1300–1308, 2017. [PubMed: 28481528]
- [4]. Nakata Maho, Shimazaki Tomomi, Hashimoto Masatomo, and Maeda Toshiyuki. Pubchemqc pm6: Data sets of 221 million molecules with optimized molecular geometries and electronic properties. *Journal of Chemical Information and Modeling*, 60(12):5891–5899, 2020. [PubMed: 33104339]
- [5]. Tox21 data browser. <https://tripod.nih.gov/tox21/>.
- [6]. Toxcast database (invitrodb). [https://epa.figshare.com/articles/dataset/ToxCast\\_Database\\_invitroDB\\_/6062623/](https://epa.figshare.com/articles/dataset/ToxCast_Database_invitroDB_/6062623/).
- [7]. Wu Zhenqin, Ramsundar Bharath, Feinberg Evan N, Gomes Joseph, Geniesse Caleb, Pappu Aneesh S, Leswing Karl, and Pande Vijay. Moleculenet: a benchmark for molecular machine learning. *Chemical science*, 9(2):513–530, 2018. [PubMed: 29629118]
- [8]. Sider side effect resource. <http://sideeffects.embl.de/>.
- [9]. Mendez David, Gaulton Anna, Bento A Patrícia, Chambers Jon, De Veij Marleen, Félix Eloy, Magariños María Paula, Mosquera Juan F, Mutowo Prudence, Nowotka Michał, et al. ChEMBL: towards direct deposition of bioassay data. *Nucleic acids research*, 47(D1):D930–D940, 2019. [PubMed: 30398643]
- [10]. Chen Lieyang, Cruz Anthony, Ramsey Steven, Dickson Callum J, Duca Jose S, Hornak Viktor, Koes David R, and Kurtzman Tom. Hidden bias in the dud-e dataset leads to misleading performance of deep learning in structure-based virtual screening. *PloS one*, 14(8):e0220113, 2019. [PubMed: 31430292]
- [11]. Keller Andreas and Vosshall Leslie B. Olfactory perception of chemically diverse molecules. *BMC neuroscience*, 17(1):1–17, 2016. [PubMed: 26728069]
- [12]. Mobley David L and Guthrie J Peter. Freesolv: a database of experimental and calculated hydration free energies, with input files. *Journal of computer-aided molecular design*, 28(7):711–720, 2014. [PubMed: 24928188]
- [13]. Delaney John S. Esol: estimating aqueous solubility directly from molecular structure. *Journal of chemical information and computer sciences*, 44(3):1000–1005, 2004. [PubMed: 15154768]
- [14]. Wang Jian-Bing, Cao Dong-Sheng, Zhu Min-Feng, Yun Yong-Huan, Xiao Nan, and Liang Yi-Zeng. In silico evaluation of log<sub>D</sub> 4 and comparison with other prediction methods. *Journal of Chemometrics*, 29(7):389–398, 2015.
- [15]. Li Wei. 10.6084/m9.figshare.6025748, Nov 2018.
- [16]. Wang Yanli, Xiao Jewen, Suzek Tugba O, Zhang Jian, Wang Jiyao, Zhou Zhigang, Han Lianyi, Karapetyan Karen, Dracheva Svetlana, Shoemaker Benjamin A, et al. Pubchem’s bioassay database. *Nucleic acids research*, 40(D1):D400–D412, 2012. [PubMed: 22140110]
- [17]. Wang Renxiao, Fang Xueliang, Lu Yipin, and Wang Shaomeng. The pdbbind database: Collection of binding affinities for protein- ligand complexes with known three-dimensional structures. *Journal of medicinal chemistry*, 47(12):2977–2980, 2004. [PubMed: 15163179]
- [18]. Liu Zhihai, Li Yan, Han Li, Li Jie, Liu Jie, Zhao Zhixiong, Nie Wei, Liu Yuchen, and Wang Renxiao. Pdb-wide collection of binding data: current status of the pdbbind database. *Bioinformatics*, 31(3):405–412, 2015. [PubMed: 25301850]
- [19]. Martins Ines Filipa, Teixeira Ana L, Pinheiro Luis, and Ic Andre O. A bayesian approach to in silico blood-brain barrier penetration modeling. *Journal of chemical information and modeling*, 52(6):1686–1697, 2012. [PubMed: 22612593]

- [20]. Goodfellow Ian, Bengio Yoshua, Courville Aaron, and Bengio Yoshua. Deep learning, volume 1. MIT press Cambridge, 2016.
- [21]. Muratov Eugene N, Bajorath Jürgen, Sheridan Robert P, Tetko Igor V, Filimonov Dmitry, Poroikov Vladimir, Oprea Tudor I, Baskin Igor I, Varnek Alexandre, Roitberg Adrian, et al. Qsar without borders. *Chemical Society Reviews*, 49(11):3525–3564, 2020. [PubMed: 32356548]
- [22]. Netzeva Tatiana I, Worth Andrew P, Aldenberg Tom, Benigni Romualdo, Cronin Mark TD, Gramatica Paola, Jaworska Joanna S, Kahn Scott, Klopman Gilles, Marchant Carol A, et al. Current status of methods for defining the applicability domain of (quantitative) structure-activity relationships: The report and recommendations of ecvam workshop 52. *Alternatives to Laboratory Animals*, 33(2):155–173, 2005. [PubMed: 16180989]
- [23]. Pan Sinno Jialin and Yang Qiang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.
- [24]. Mysinger Michael M, Carchia Michael, Irwin John J, and Shoichet Brian K. Directory of useful decoys, enhanced (dud-e): better ligands and decoys for better benchmarking. *Journal of medicinal chemistry*, 55(14):6582–6594, 2012. [PubMed: 22716043]
- [25]. Glavatskikh Marta, Leguy Jules, Hunault Gilles, Cauchy Thomas, and Da Mota Benoit. Dataset’s chemical diversity limits the generalizability of machine learning predictions. *Journal of Cheminformatics*, 11(1):1–15, 2019. [PubMed: 30604073]
- [26]. Cai Chenjing, Wang Shiwei, Xu Youjun, Zhang Weilin, Tang Ke, Ouyang Qi, Lai Luhua, and Pei Jianfeng. Transfer learning for drug discovery. *Journal of Medicinal Chemistry*, 63(16):8683–8694, 2020. [PubMed: 32672961]
- [27]. Li Xinhao and Fourches Denis. Inductive transfer learning for molecular activity prediction: Next-gen qsar models with molmpofit. *Journal of Cheminformatics*, 12:1–15, 2020. [PubMed: 33430988]
- [28]. Goh Garrett B, Siegel Charles, Vishnu Abhinav, and Hodas Nathan. Using rule-based labels for weak supervised learning: a chemnet for transferable chemical property prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 302–310, 2018.
- [29]. Hüllermeier Eyke and Waegeman Willem. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *arXiv preprint arXiv:1910.09457*, 2019.
- [30]. Breiman Leo. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [31]. Sheridan Robert P. Three useful dimensions for domain applicability in qsar models using random forest. *Journal of chemical information and modeling*, 52(3):814–823, 2012. [PubMed: 22385389]
- [32]. Toplak Marko, Mo nik Rok, Polajnar Matija, Bosni Zoran, Carlsson Lars, Hasselgren Catrin, Demšar Janez, Boyer Scott, Zupan Blaž, and Stålring Jonna. Assessment of machine learning reliability methods for quantifying the applicability domain of qsar regression models. *Journal of chemical information and modeling*, 54(2):431–441, 2014. [PubMed: 24490838]
- [33]. Lakshminarayanan Balaji, Pritzel Alexander, and Blundell Charles. Simple and scalable predictive uncertainty estimation using deep ensembles. *arXiv preprint arXiv:1612.01474*, 2016.
- [34]. Scalia Gabriele, Grambow Colin A, Pernici Barbara, Li Yi-Pei, and Green William H. Evaluating scalable uncertainty estimation methods for deep learning-based molecular property prediction. *Journal of chemical information and modeling*, 60(6):2697–2717, 2020. [PubMed: 32243154]
- [35]. Nix David A and Weigend Andreas S. Estimating the mean and variance of the target probability distribution. In *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN’94)*, volume 1, pages 55–60. IEEE, 1994.
- [36]. Svensson Fredrik, Norinder Ulf, and Bender Andreas. Modelling compound cytotoxicity using conformal prediction and pubchem hts data. *Toxicology research*, 6(1):73–80, 2017. [PubMed: 30090478]
- [37]. Norinder Ulf, Carlsson Lars, Boyer Scott, and Eklund Martin. Introducing conformal prediction in predictive modeling. a transparent and flexible alternative to applicability domain determination. *Journal of chemical information and modeling*, 54(6):1596–1603, 2014. [PubMed: 24797111]

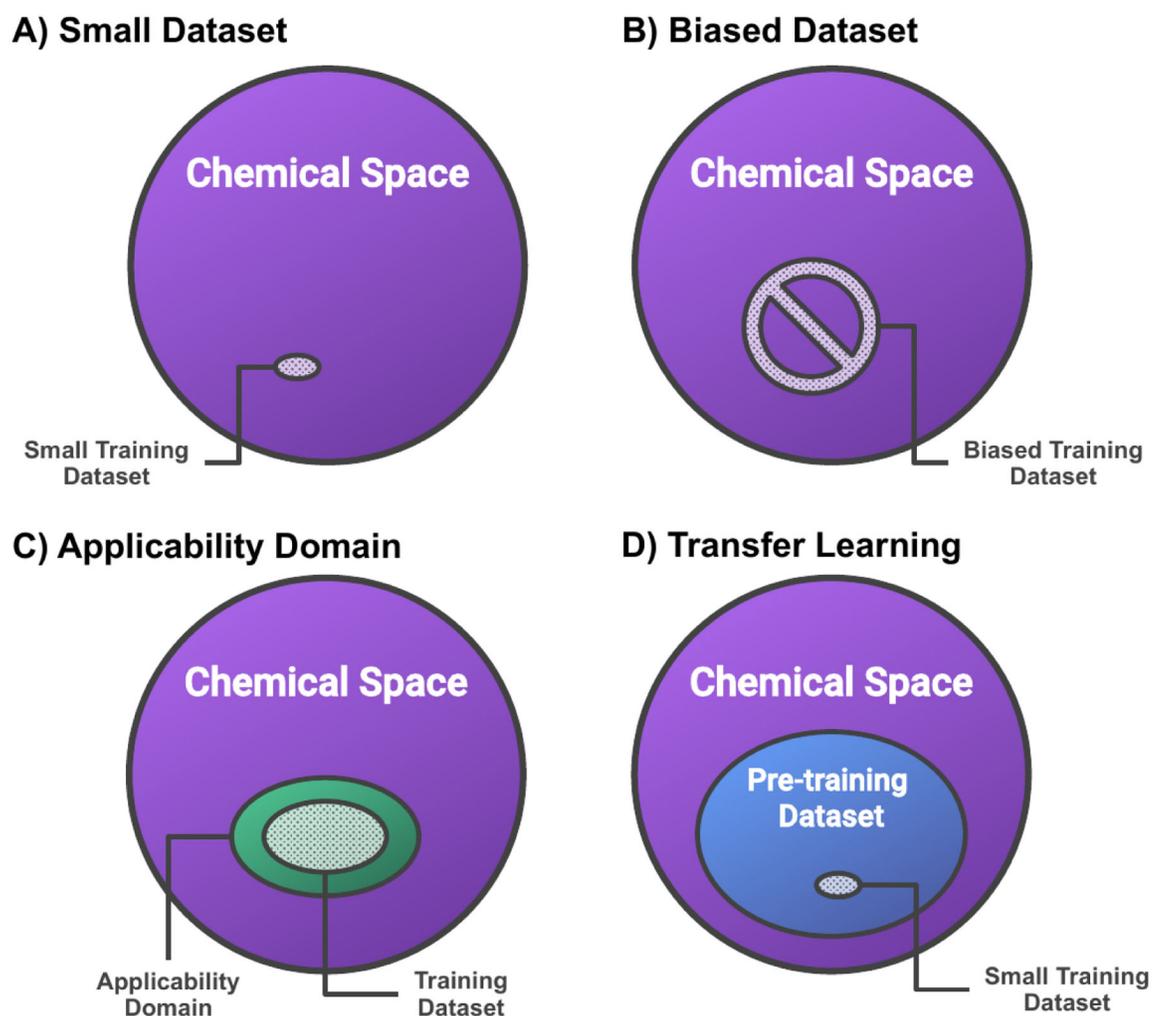
- [38]. Svensson Fredrik, Afzal Avid M, Norinder Ulf, and Bender Andreas. Maximizing gain in high-throughput screening using conformal prediction. *Journal of cheminformatics*, 10(1):1–10, 2018. [PubMed: 29340790]
- [39]. Cortés-Ciriano Isidro and Bender Andreas. Concepts and applications of conformal prediction in computational drug discovery. *arXiv preprint arXiv:1908.03569*, 2019.
- [40]. Rasmussen Carl Edward. Gaussian processes in machine learning. In *Summer school on machine learning*, pages 63–71. Springer, 2003.
- [41]. Sahli Costabal Francisco, Matsuno Kristen, Yao Jiang, Perdikaris Paris, and Kuhl Ellen. Machine learning in drug development: Characterizing the effect of 30 drugs on the qt interval using gaussian process regression, sensitivity analysis, and uncertainty quantification. *Computer Methods in Applied Mechanics and Engineering*, 348:313–333, 2019. [PubMed: 32863454]
- [42]. Bannan Caitlin C, Mobley David L, and Skillman A Geoffrey. Sampl6 challenge results from pka predictions based on a general gaussian process model. *Journal of computer-aided molecular design*, 32(10):1165–1177, 2018. [PubMed: 30324305]
- [43]. Hie Brian, Bryson Bryan D, and Berger Bonnie. Leveraging uncertainty in machine learning accelerates biological discovery and design. *Cell Systems*, 11(5):461–477, 2020. [PubMed: 33065027]
- [44]. Blundell Charles, Cornebise Julien, Kavukcuoglu Koray, and Wierstra Daan. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pages 1613–1622. PMLR, 2015.
- [45]. Zhang Yao et al. Bayesian semi-supervised learning for uncertainty-calibrated prediction of molecular properties and active learning. *Chemical science*, 10(35):8154–8163, 2019. [PubMed: 31857882]
- [46]. Guo Chuan, Pleiss Geoff, Sun Yu, and Weinberger Kilian Q. On calibration of modern neural networks. In *International Conference on Machine Learning*, pages 1321–1330. PMLR, 2017.
- [47]. Williams Dominic P, Lazic Stanley E, Foster Alison J, Semenova Elizaveta, and Morgan Paul. Predicting drug-induced liver injury with bayesian machine learning. *Chemical research in toxicology*, 33(1):239–248, 2019. [PubMed: 31535850]
- [48]. Semenova Elizaveta, Williams Dominic P, Afzal Avid M, and Lazic Stanley E. A bayesian neural network for toxicity prediction. *Computational Toxicology*, 16:100133, 2020.
- [49]. Lazic Stanley E, Edmunds Nicholas, and Pollard Christopher E. Predicting drug safety and communicating risk: benefits of a bayesian approach. *Toxicological Sciences*, 162(1):89–98, 2018. [PubMed: 29126124]
- [50]. Settles Burr. *Active learning literature survey*. 2009.
- [51]. Reker Daniel and Schneider Gisbert. Active-learning strategies in computer-assisted drug discovery. *Drug discovery today*, 20(4):458–465, 2015. [PubMed: 25499665]
- [52]. Reker Daniel, Schneider Petra, Schneider Gisbert, and Brown JB. Active learning for computational chemogenomics. *Future medicinal chemistry*, 9(4):381–402, 2017. [PubMed: 28263088]
- [53]. Ryu Seongok, Kwon Yongchan, and Kim Woo Youn. A bayesian graph convolutional network for reliable prediction of molecular properties with uncertainty quantification. *Chemical science*, 10(36):8438–8446, 2019. [PubMed: 31803423]
- [54]. Hirschfeld Lior, Swanson Kyle, Yang Kevin, Barzilay Regina, and Coley Connor W. Uncertainty quantification using neural networks for molecular property prediction. *Journal of Chemical Information and Modeling*, 60(8):3770–3780, 2020. [PubMed: 32702986]
- [55]. Moss Henry B and Griffiths Ryan-Rhys. Gaussian process molecule property prediction with flowmo. *arXiv preprint arXiv:2010.01118*, 2020.
- [56]. Miteva Maria A, Guyon Frederic, and Tufféry Pierre. Frog2: Efficient 3d conformation ensemble generator for small compounds. *Nucleic acids research*, 38:W622–W627, 2010. [PubMed: 20444874]
- [57]. Guasch Laura, Peach Megan L, and Nicklaus Marc C. Tautomerism of warfarin: combined cheminformatics, quantum chemical, and nmr investigation. *The Journal of organic chemistry*, 80(20):9900–9909, 2015. [PubMed: 26372257]

- [58]. Weininger David. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1):31–36, 1988.
- [59]. Girard Agathe, Rasmussen Carl Edward, Quinonero-Candela Joaquin, and Murray-Smith Roderick. Gaussian process priors with uncertain inputs? application to multiple-step ahead time series forecasting. 2003.
- [60]. Hanafusa Ryo and Okadome Takeshi. Bayesian kernel regression for noisy inputs based on nadaraya–watson estimator constructed from noiseless training data. *Advances in Data Science and Adaptive Analysis*, 12(01):2050004, 2020.
- [61]. Titsias Michalis and Lawrence Neil D. Bayesian gaussian process latent variable model. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 844–851. *JMLR Workshop and Conference Proceedings*, 2010.
- [62]. Li Ping and Chen Songcan. A review on gaussian process latent variable models. *CAAI Transactions on Intelligence Technology*, 1(4):366–376, 2016.
- [63]. Axelrod Simon and Gomez-Bombarelli Rafael. Molecular machine learning with conformer ensembles. arXiv preprint arXiv:2012.08452, 2020.
- [64]. Trott Oleg and Olson Arthur J. Autodock vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *Journal of computational chemistry*, 31(2):455–461, 2010. [PubMed: 19499576]
- [65]. Friesner Richard A, Banks Jay L, Murphy Robert B, Halgren Thomas A, Klicic Jasna J, Mainz Daniel T, Repasky Matthew P, Knoll Eric H, Shelley Mee, Perry Jason K, et al. Glide: a new approach for rapid, accurate docking and scoring. 1. method and assessment of docking accuracy. *Journal of medicinal chemistry*, 47(7):1739–1749, 2004. [PubMed: 15027865]
- [66]. Spitzer Russell and Jain Ajay N. Surfex-dock: Docking benchmarks and real-world application. *Journal of computer-aided molecular design*, 26(6):687–699, 2012. [PubMed: 22569590]
- [67]. Verdonk Marcel L, Cole Jason C, Hartshorn Michael J, Murray Christopher W, and Taylor Richard D. Improved protein–ligand docking using gold. *Proteins: Structure, Function, and Bioinformatics*, 52(4):609–623, 2003.
- [68]. Cheng Tiejun, Li Xun, Li Yan, Liu Zhihai, and Wang Renxiao. Comparative assessment of scoring functions on a diverse test set. *Journal of chemical information and modeling*, 49(4):1079–1093, 2009. [PubMed: 19358517]
- [69]. Li Yan, Han Li, Liu Zhihai, and Wang Renxiao. Comparative assessment of scoring functions on an updated benchmark: 2. evaluation methods and general results. *Journal of chemical information and modeling*, 54(6):1717–1736, 2014. [PubMed: 24708446]
- [70]. Khamis Mohamed A, Gomaa Walid, and Ahmed Walaa F. Machine learning in computational docking. *Artificial intelligence in medicine*, 63(3):135–152, 2015. [PubMed: 25724101]
- [71]. Procacci Piero. Methodological uncertainties in drug-receptor binding free energy predictions based on classical molecular dynamics. *Current Opinion in Structural Biology*, 67:127–134, 2021. [PubMed: 33220532]
- [72]. Jiménez José, Skalic Miha, Martínez-Rosell Gerard, and De Fabritiis Gianni. K deep: protein–ligand absolute binding affinity prediction via 3d-convolutional neural networks. *Journal of chemical information and modeling*, 58(2):287–296, 2018. [PubMed: 29309725]
- [73]. Pansar Tatu and Poso Antti. Binding affinity via docking: fact and fiction. *Molecules*, 23(8):1899, 2018.
- [74]. Morrone Joseph A, Weber Jeffrey K, Huynh Tien, Luo Heng, and Cornell Wendy D. Combining docking pose rank and structure with deep learning improves protein–ligand binding mode prediction over a baseline docking approach. *Journal of chemical information and modeling*, 60(9):4170–4179, 2020. [PubMed: 32077698]
- [75]. Shirts Michael R, Mobley David L, and Chodera John D. Alchemical free energy calculations: ready for prime time? *Annual reports in computational chemistry*, 3:41–59, 2007.
- [76]. Shirts Michael R and Mobley David L. An introduction to best practices in free energy calculations. In *Biomolecular Simulations*, pages 271–311. Springer, 2013.

- [77]. Gapsys Vytautas, Michielssens Servaas, Peters Jan Henning, de Groot Bert L, and Leonov Hadas. Calculation of binding free energies. In *Molecular Modeling of Proteins*, pages 173–209. Springer, 2015.
- [78]. Cournia Zoe, Allen Bryce K, Beuming Thijs, Pearlman David A, Radak Brian K, and Sherman Woody. Rigorous free energy simulations in virtual screening. *Journal of Chemical Information and Modeling*, 60(9):4153–4169, 2020. [PubMed: 32539386]
- [79]. Wang Lingle, Wu Yujie, Deng Yuqing, Kim Byungchan, Pierce Levi, Krilov Goran, Lupyan Dmitry, Robinson Shaughnessy, Dahlgren Markus K, Greenwood Jeremy, et al. Accurate and reliable prediction of relative ligand binding potency in prospective drug discovery by way of a modern free-energy calculation protocol and force field. *Journal of the American Chemical Society*, 137(7):2695–2703, 2015. [PubMed: 25625324]
- [80]. Aldeghi Matteo, Heifetz Alexander, Bodkin Michael J, Knapp Stefan, and Biggin Philip C. Accurate calculation of the absolute free energy of binding for drug molecules. *Chemical science*, 7(1):207–218, 2016. [PubMed: 26798447]
- [81]. Shan Yibing, Kim Eric T, Eastwood Michael P, Dror Ron O, Seeliger Markus A, and Shaw David E. How does a drug molecule find its target binding site? *Journal of the American Chemical Society*, 133(24):9181–9183, 2011. [PubMed: 21545110]
- [82]. Zwanzig Robert W. High-temperature equation of state by a perturbation method. i. nonpolar gases. *The Journal of Chemical Physics*, 22(8):1420–1426, 1954.
- [83]. Kirkwood John G. Statistical mechanics of fluid mixtures. *The Journal of chemical physics*, 3(5):300–313, 1935.
- [84]. Kumar Shankar, Rosenberg John M, Bouzida Djamel, Swendsen Robert H, and Kollman Peter A. The weighted histogram analysis method for free-energy calculations on biomolecules. i. the method. *Journal of computational chemistry*, 13(8):1011–1021, 1992.
- [85]. Bennett Charles H. Efficient estimation of free energy differences from monte carlo data. *Journal of Computational Physics*, 22(2):245–268, 1976.
- [86]. Shirts Michael R and Chodera John D. Statistically optimal analysis of samples from multiple equilibrium states. *The Journal of chemical physics*, 129(12):124105, 2008. [PubMed: 19045004]
- [87]. Wu Hao, Paul Fabian, Wehmeyer Christoph, and Noé Frank. Multiensemble markov models of molecular thermodynamics and kinetics. *Proceedings of the National Academy of Sciences*, 113(23):E3221–E3230, 2016.
- [88]. Rosta Edina and Hummer Gerhard. Free energies from dynamic weighted histogram analysis using unbiased markov state model. *Journal of chemical theory and computation*, 11(1):276–285, 2015. [PubMed: 26574225]
- [89]. Grossfield Alan, Patrone Paul N, Roe Daniel R, Schultz Andrew J, Siderius Daniel W, and Zuckerman Daniel M. Best practices for quantification of uncertainty and sampling quality in molecular simulations [article v1. 0]. *Living journal of computational molecular science*, 1(1), 2018.
- [90]. Chodera John D. A simple method for automated equilibration detection in molecular simulations. *Journal of chemical theory and computation*, 12(4):1799–1805, 2016. [PubMed: 26771390]
- [91]. Crooks Gavin E. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Physical Review E*, 60(3):2721, 1999.
- [92]. Gapsys Vytautas, Pérez-Benito Laura, Aldeghi Matteo, Seeliger Daniel, Van Vlijmen Herman, Tresadern Gary, and de Groot Bert L. Large scale relative protein ligand binding affinities using non-equilibrium alchemy. *Chemical Science*, 11(4):1140–1152, 2020.
- [93]. Baumann Hannah, Gapsys Vytautas, de Groot Bert L, and Mobley David. Challenges encountered applying equilibrium and non-equilibrium binding free energy calculations. 2020.
- [94]. Khalak Yuri, Tresadern Gary, de Groot Bert L, and Gapsys Vytautas. Non-equilibrium approach for binding free energies in cyclodextrins in sampl7: force fields and software. *Journal of computer-aided molecular design*, pages 1–13, 2020. [PubMed: 31792884]
- [95]. Gapsys Vytautas and de Groot Bert L. On the importance of statistics in molecular simulations for thermodynamics, kinetics and simulation box size. *Elife*, 9:e57589, 2020. [PubMed: 32812868]

- [96]. Shirts Michael R, Klein Christoph, Swails Jason M, Yin Jian, Gilson Michael K, Mobley David L, Case David A, and Zhong Ellen D. Lessons learned from comparing molecular dynamics engines on the *samp15* dataset. *Journal of computer-aided molecular design*, 31(1):147–161, 2017. [PubMed: 27787702]
- [97]. Merz Pascal T and Shirts Michael R. Testing for physical validity in molecular simulations. *PLoS one*, 13(9):e0202764, 2018. [PubMed: 30188933]
- [98]. Beauchamp Kyle A, Lin Yu-Shan, Das Rhiju, and Pande Vijay S. Are protein force fields getting better? a systematic benchmark on 524 diverse nmr measurements. *Journal of chemical theory and computation*, 8(4):1409–1414, 2012. [PubMed: 22754404]
- [99]. Nerenberg Paul S and Head-Gordon Teresa. New developments in force fields for biomolecular simulations. *Current opinion in structural biology*, 49:129–138, 2018. [PubMed: 29477047]
- [100]. Mobley David L, Bannan Caitlin C, Rizzi Andrea, Bayly Christopher I, Chodera John D, Lim Victoria T, Lim Nathan M, Beauchamp Kyle A, Slochower David R, Shirts Michael R, et al. Escaping atom types in force fields using direct chemical perception. *Journal of chemical theory and computation*, 14(11):6076–6092, 2018. [PubMed: 30351006]
- [101]. Wang Lee-Ping, Martinez Todd J, and Pande Vijay S. Building force fields: An automatic, systematic, and reproducible approach. *The journal of physical chemistry letters*, 5(11):1885–1891, 2014. [PubMed: 26273869]
- [102]. Wang Lee-Ping, McKiernan Keri A, Gomes Joseph, Beauchamp Kyle A, Head-Gordon Teresa, Rice Julia E, Swope William C, Martínez Todd J, and Pande Vijay S. Building a more predictive protein force field: a systematic and reproducible route to *amber-fb15*. *The Journal of Physical Chemistry B*, 121(16):4023–4039, 2017. [PubMed: 28306259]
- [103]. Qiu Yudong, Nerenberg Paul S., Head-Gordon Teresa, and Wang Lee-Ping. Systematic optimization of water models using liquid/vapor surface tension data. *The Journal of Physical Chemistry B*, 123(32):7061–7073, 2019. [PubMed: 31314516]
- [104]. Ponder Jay W, Wu Chuanjie, Ren Pengyu, Pande Vijay S, Chodera John D, Schnieders Michael J, Haque Imran, Mobley David L, Lambrecht Daniel S, DiStasio Robert A Jr, et al. Current status of the amoeba polarizable force field. *The journal of physical chemistry B*, 114(8):2549–2564, 2010. [PubMed: 20136072]
- [105]. Shi Yuanjun, Laury Marie L, Wang Zhi, and Ponder Jay W. Amoeba binding free energies for the *samp17* trimertrip host–guest challenge. *Journal of Computer-Aided Molecular Design*, pages 1–15, 2020. [PubMed: 31792884]
- [106]. Smith Justin S, Isayev Olexandr, and Roitberg Adrian E. Ani-1: an extensible neural network potential with dft accuracy at force field computational cost. *Chemical science*, 8(4):3192–3203, 2017. [PubMed: 28507695]
- [107]. Lahey Shae-Lynn J and Rowley Christopher N. Simulating protein–ligand binding with neural network potentials. *Chemical Science*, 11(9):2362–2368, 2020. [PubMed: 34084397]
- [108]. I ik Mehtap, Bergazin Teresa Danielle, Fox Thomas, Rizzi Andrea, Chodera John D, and Mobley David L. Assessing the accuracy of octanol–water partition coefficient predictions in the *samp16* part ii log p challenge. *Journal of computer-aided molecular design*, 34(4):335–370, 2020. [PubMed: 32107702]
- [109]. Sanchez-Lengeling Benjamin and Aspuru-Guzik Alán. Inverse molecular design using machine learning: Generative models for matter engineering. *Science*, 361(6400):360–365, 2018. [PubMed: 30049875]
- [110]. Pollice Robert, dos Passos Gomes Gabriel, Aldeghi Matteo, Hickman Riley J, Krenn Mario, Lavigne Cyrille, Lindner-D’Addario Michael, Nigam AkshatKumar, Ser Cher Tian, Yao Zhenpeng, et al. Data-driven strategies for accelerated materials design. *Accounts of Chemical Research*, pages 1120–1127, 2021.
- [111]. Kingma Diederik P and Welling Max. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [112]. Goodfellow Ian, Pouget-Abadie Jean, Mirza Mehdi, Xu Bing, Warde-Farley David, Ozair Sherjil, Courville Aaron, and Bengio Yoshua. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

- [113]. Sutton Richard S and Barto Andrew G. Reinforcement learning: An introduction. MIT press, 2018.
- [114]. Mitchell Melanie. An introduction to genetic algorithms. MIT press, 1998.
- [115]. Gómez-Bombarelli Rafael, Wei Jennifer N, Duvenaud David, Hernández-Lobato José Miguel, Sánchez-Lengeling Benjamín, Sheberla Dennis, AguileraIparraquirre Jorge, Hirzel Timothy D, Adams Ryan P, and Aspuru-Guzik Alán. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2):268–276, 2018. [PubMed: 29532027]
- [116]. Guimaraes Gabriel Lima, Sanchez-Lengeling Benjamin, Outeiral Carlos, Farias Pedro Luis Cunha, and Aspuru-Guzik Alán. Objective-reinforced generative adversarial networks (organ) for sequence generation models. *arXiv preprint arXiv:1705.10843*, 2017.
- [117]. Sanchez-Lengeling Benjamin, Outeiral Carlos, Guimaraes Gabriel L, and Aspuru-Guzik Alan. Optimizing distributions over molecular space. an objective-reinforced generative adversarial network for inverse-design chemistry (organic). *ChemRxiv*, 2017, 2017.
- [118]. Ståhl Niclas, Falkman Goran, Karlsson Alexander, Mathiason Gunnar, and Bostrom Jonas. Deep reinforcement learning for multiparameter optimization in de novo drug design. *Journal of chemical information and modeling*, 59(7):3166–3176, 2019. [PubMed: 31273995]
- [119]. Zhou Zhenpeng, Kearnes Steven, Li Li, Zare Richard N, and Riley Patrick. Optimization of molecules via deep reinforcement learning. *Scientific reports*, 9(1):1–10, 2019. [PubMed: 30626917]
- [120]. Popova Mariya, Isayev Olexandr, and Tropsha Alexander. Deep reinforcement learning for de novo drug design. *Science advances*, 4(7):eaap7885, 2018. [PubMed: 30050984]
- [121]. Nigam AkshatKumar, Friederich Pascal, Krenn Mario, and Aspuru-Guzik. Augmenting genetic algorithms with deep neural networks for exploring the chemical space. *arXiv preprint arXiv:1909.11655*, 2019.
- [122]. Jensen Jan H. A graph-based genetic algorithm and generative model/monte carlo tree search for the exploration of chemical space. *Chemical science*, 10(12):3567–3572, 2019. [PubMed: 30996948]
- [123]. Yoshikawa Naruki, Terayama Kei, Sumita Masato, Homma Teruki, Oono Kenta, and Tsuda Koji. Population-based de novo molecule generation, using grammatical evolution. *Chemistry Letters*, 47(11):1431–1434, 2018.
- [124]. Thiede Luca A, Krenn Mario, Nigam AkshatKumar, and Aspuru-Guzik Alan. Curiosity in exploring chemical space: Intrinsic rewards for deep molecular reinforcement learning. *arXiv preprint arXiv:2012.11293*, 2020.
- [125]. Shen Cynthia, Krenn Mario, Eppel Sagi, and Aspuru-Guzik Alan. Deep molecular dreaming: Inverse machine learning for de-novo molecular design and interpretability with surjective representations. *arXiv preprint arXiv:2012.09712*, 2020.
- [126]. Krenn Mario, Häse Florian, Nigam AkshatKumar, Friederich Pascal, and Aspuru-Guzik Alan. Self-referencing embedded strings (selfies): A 100% robust molecular string representation. *Machine Learning: Science and Technology*, 1(4):045024, 2020.
- [127]. Nigam A, Pollice Robert, Krenn Mario, dos Passos Gomes Gabriel, and Aspuru-Guzik Alan. Beyond generative models: Superfast traversal, optimization, novelty, exploration and discovery (stoned) algorithm for molecules using selfies. 2020.



**FIG. 1.** Common problems with datasets for data-driven molecular property prediction and strategies to improve confidence in the models. A) Small training datasets hamper generalization. B) Biased training datasets lead to models only learning the inherent bias rather than meaningful relationships. C) The applicability domain of a model defines the molecules for which the property predictions are expected to be reliable. D) Transfer learning via pre-training on large (unlabelled) datasets improves prediction accuracy for property prediction based on small training datasets.

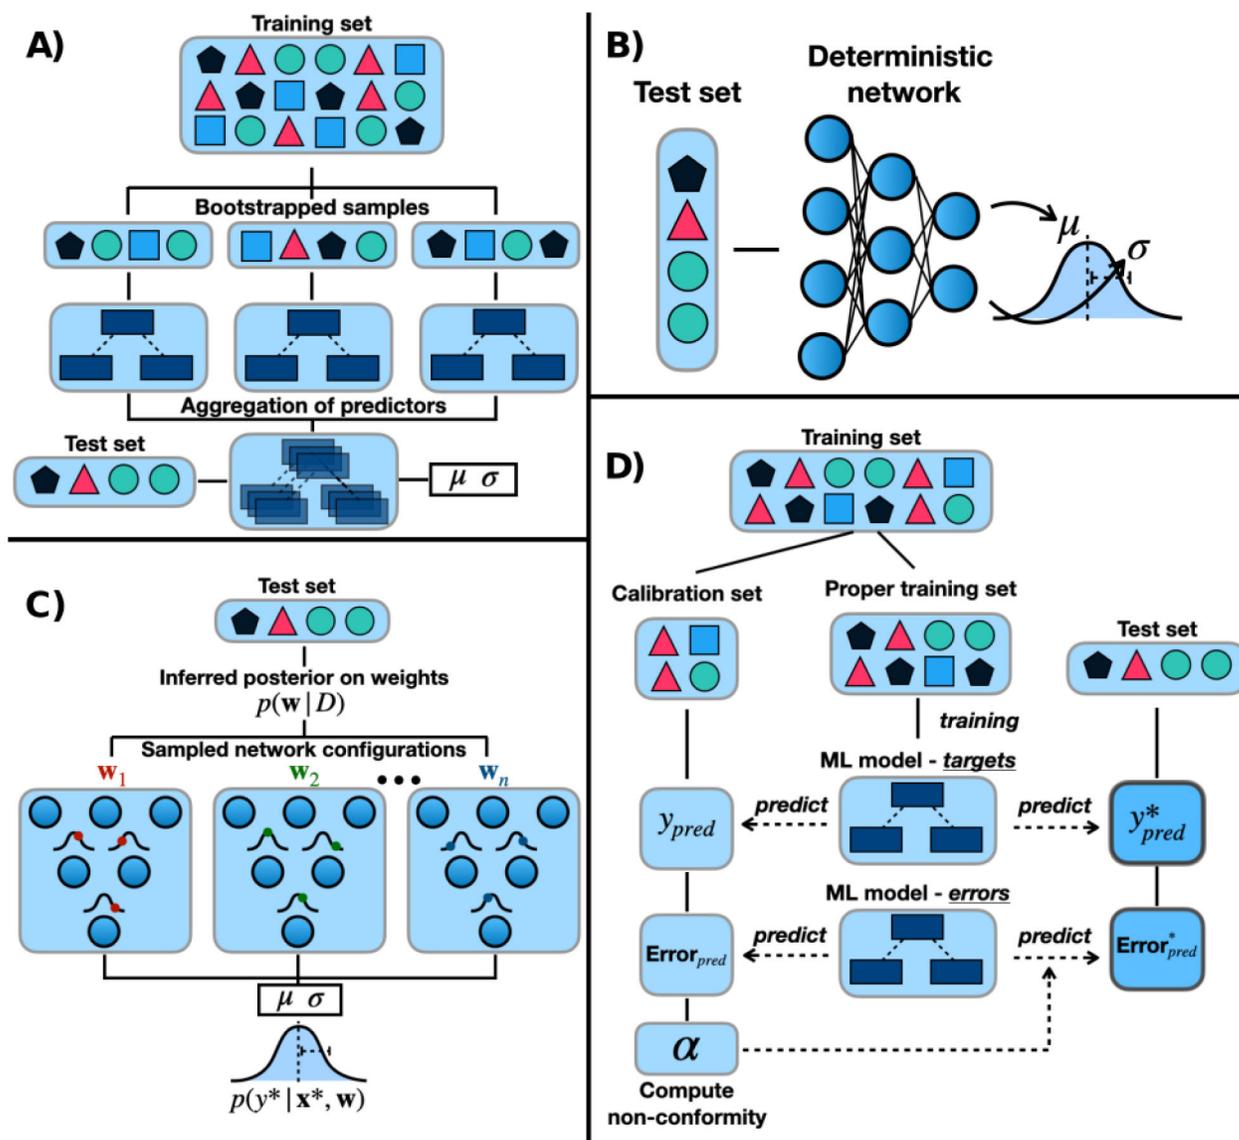
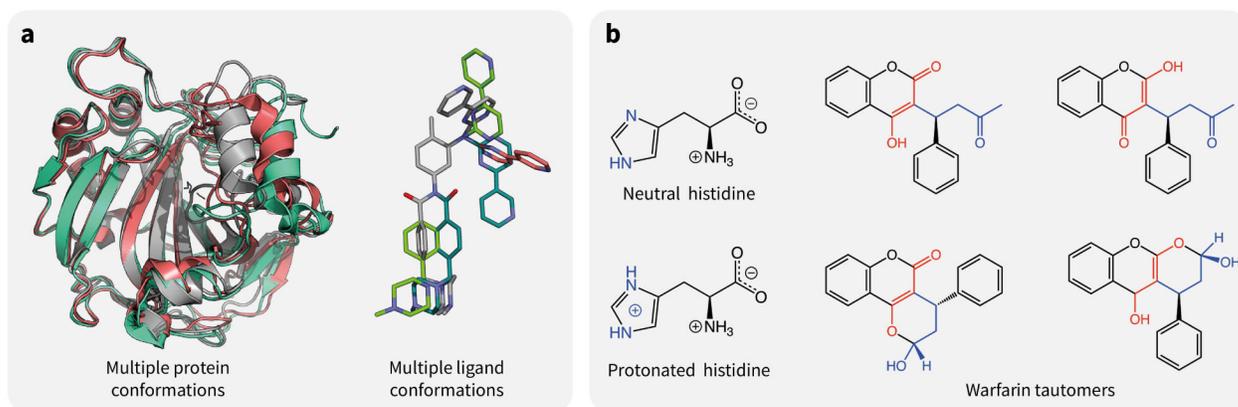


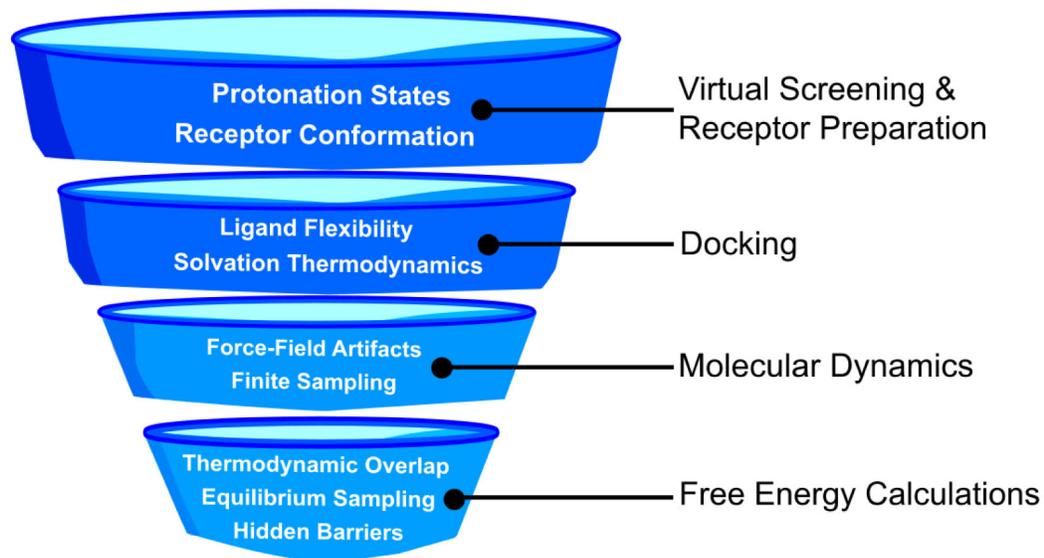
FIG. 2.

Four common methods for producing uncertainties or confidence intervals in molecular property prediction. A) Bagging (bootstrap aggregation) algorithm with random forests as the ML model. B) Deterministic neural network trained using frequentist maximum likelihood estimation to approximate the distribution on the targets. C) Bayesian neural network in a regression setting depicting the test time sampling of different network configurations from the posterior distribution on the weights. D) Inductive conformal prediction in a regression setting using random forests as the ML model.

**FIG. 3.**

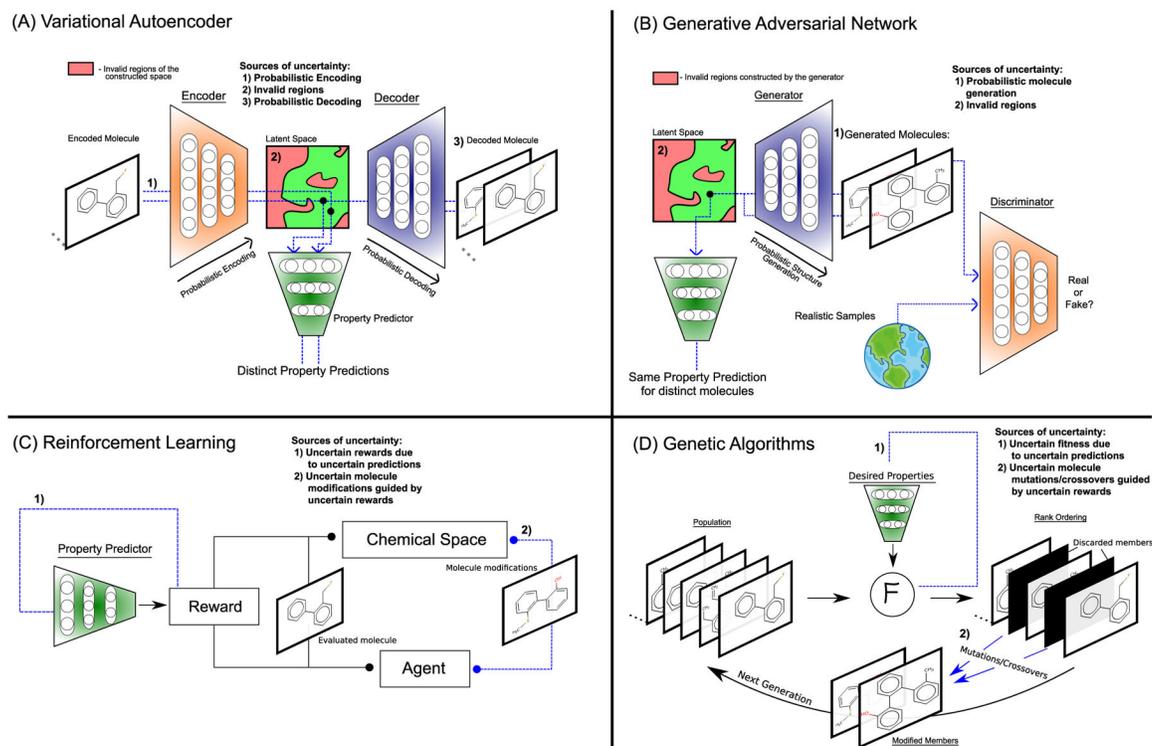
Examples of uncertainty in the structural and chemical inputs used by ML models. (a) Conformational variability is encountered both in proteins and small organic molecules. On the left-hand side, some of the conformations of carbonic anhydrase II are shown, as observed via nuclear magnetic resonance (PDB-ID 6HD2). On the right-hand side are some of the possible conformations of Imatinib (computed with Frog [56]). (b) Examples of major protonation and tautomeric states present in the solution for an amino acid and a drug. The protonated form of histidine has a  $pK_a$  of approximately 6.0, which means that it is generally expected to be found in its neutral form at physiological pH. Yet, because of the local environment in protein binding pockets, it is often unclear which of the two is the major species, and they may both co-exist in non-negligible proportions. On the right-hand side, some of the tautomeric states of warfarin [57] are shown.

## Sources of Uncertainty Arising from Simulated Binding Affinity Calculations



**FIG. 4.**

A generalized sequence of steps to compute binding affinities, and their corresponding facets of uncertainty. The top half of the diagram illustrates low-cost methods, where the sources of uncertainty can be more easily identified. The bottom half of the pathway refers to the methods which sacrifice computational cost for accuracy, and for which the minimization of uncertainty largely represents open problems in computing accurate binding affinities.



**FIG. 5.** Architectures and sources of uncertainty of generative models for molecular design, namely: (A) Variational autoencoder, (B) Generative adversarial network, (C) Reinforcement learning, and (D) Genetic algorithms. Sources of uncertainty within each architecture are highlighted in blue.

**TABLE I.**  
Some of the common datasets used for molecular machine learning and data-driven property prediction.

Name	Description	Number of molecules	Possible bias
ZINC [1]	Database of commercially available compounds together with very simple estimated molecular properties for virtual screening.	1.4 billion	Inherently biased by currently synthesizable chemical space. Consequently, the molecular shapes have been shown to be highly biased against sphere-like molecules.
QM9 [2]	Electronic properties estimated using density functional theory (DFT) simulations.	134 thousand	Biased towards small molecules only containing the elements C, H, N, O and F.
PubChemQC [3, 4]	Geometries and electronic properties of molecules with short string representations taken from PubChem.	221 million	Biased towards small molecules that have been reported in the literature before.
Tox21 [5]	Toxicologic properties of molecules with respect to 12 different assays	13 thousand	Biased towards environmental compounds and approved drugs.
ToxCast [6]	High-throughput screening and computational data for the toxicology of molecules from industry, consumer products and the food industry based on cell assays.	1.8 thousand	Biased towards molecules used in industry, consumer products and the food industry.
ClinTox [7]	Drugs and drug candidates that made it to clinical trials and were either approved or failed.	1.5 thousand	Biased towards drugs that made it to clinical trials.
SIDER [8]	Recorded adverse drug reactions of marketed drugs.	1.4 thousand	Biased towards marketed drugs.
ChEMBL [9]	Bioactive small molecules and their activities extracted from the literature, from clinical trials and from other databases.	2.0 million	Biased towards compounds for which bioactivity was published in the scientific literature.
DUD-E [10]	Ligand binding affinities against 102 distinct target proteins with both strong and weak binders.	23 thousand	Biased towards molecules that have been synthesized and evaluated for binding affinity.
AqSolDB [11]	Aqueous solubility data of organic molecules taken from 9 different datasets.	10 thousand	Biased towards organic molecules with relatively high aqueous solubility.
Olfaction Prediction Challenge [11]	Olfactory perception of organic molecules at different concentrations.	0.5 thousand	Biased towards small and volatile organic molecules. Results biased by familiarity of smells.
FreeSolv [12]	Experimental and computed hydration free energies of small and neutral molecules.	0.6 thousand	Biased towards small and neutral molecules that have been studied in the literature both computationally and experimentally for hydration free energies.
ESOL [13]	Experimental aqueous solubility combining datasets for small molecules from the literature, for medium-sized molecules used as pesticides and larger proprietary compounds from the pharmaceutical industry.	2.9 thousand	The sub-groups each have a different bias as they each have different application domains.
Lipophilicity [14, 15]	Experimental n-octanol/water (buffered at pH 7.4) distribution coefficient of organic molecules taken from other databases.	1.1 thousand	Biased towards molecules with distribution coefficients between -10 and 10.
PubChem Bioassay [16]	Bioactivity outcomes from high-throughput screenings of molecules.	2.3 million	Biased towards molecules of interest and molecules that are synthesizable.
PDBbind [17, 18]	Experimental binding affinity for biomolecular complexes deposited in the protein data bank (PDB).	21.4 thousand	Biased towards complexes with available crystal structures.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Name	Description	Number of molecules	Possible bias
BBBP [19]	The blood-brain penetration partition coefficient for molecules collected from the literature.	2.1 thousand	Biased towards molecules studied in the literature for blood-brain penetration.