

Review

Computer-aided multi-objective optimization in small molecule discovery

Jenna C. Fromer¹ and Connor W. Coley^{1,2,*}¹Department of Chemical Engineering, MIT, Cambridge, MA 02139, USA²Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA 02139, USA*Correspondence: ccoley@mit.edu<https://doi.org/10.1016/j.patter.2023.100678>

THE BIGGER PICTURE Data-driven approaches to molecular design promise to accelerate the discovery of novel pharmaceuticals, electronic materials, and other small molecules. As performance of computer-aided molecular optimization for single-objective tasks has improved significantly in recent years, emerging efforts are increasingly emphasizing multi-objective design.

This review details the use of Pareto optimization algorithms for multi-objective optimization of molecular properties. Pareto optimization introduces additional algorithmic complexities, but reveals more information about the trade-offs between objectives and is more robust than scalarization approaches for this reason. There is an opportunity to adopt model-guided optimization techniques into the design of new molecules with far-reaching applications including and beyond drug discovery.



Development/Pre-production: Data science output has been rolled out/validated across multiple domains/problems

SUMMARY

Molecular discovery is a multi-objective optimization problem that requires identifying a molecule or set of molecules that balance multiple, often competing, properties. Multi-objective molecular design is commonly addressed by combining properties of interest into a single objective function using scalarization, which imposes assumptions about relative importance and uncovers little about the trade-offs between objectives. In contrast to scalarization, Pareto optimization does not require knowledge of relative importance and reveals the trade-offs between objectives. However, it introduces additional considerations in algorithm design. In this review, we describe pool-based and *de novo* generative approaches to multi-objective molecular discovery with a focus on Pareto optimization algorithms. We show how pool-based molecular discovery is a relatively direct extension of multi-objective Bayesian optimization and how the plethora of different generative models extend from single-objective to multi-objective optimization in similar ways using non-dominated sorting in the reward function (reinforcement learning) or to select molecules for retraining (distribution learning) or propagation (genetic algorithms). Finally, we discuss some remaining challenges and opportunities in the field, emphasizing the opportunity to adopt Bayesian optimization techniques into multi-objective *de novo* design.

INTRODUCTION

Molecular discovery is inherently a constrained multi-objective optimization problem. Almost every molecular design application requires multiple properties to be optimized or constrained. For example, for a new drug to be successful, it must simultaneously be potent, bioavailable, safe, and synthesizable. Multi-objective optimization, also referred to as multi-parameter optimization, pertains to other applications as well, including solvent design,^{1–4} personal care products,^{5,6} electronic materials,^{7–11} functional polymers,^{12,13} and other materials.^{14–16} Redox-active species in redox flow batteries must maximize redox potential and solubi-

lity to ensure a high cell voltage.^{17,18} Sustainability of new materials (e.g., emissions caused during production and disposal¹⁹) is also an increasingly important design objective,^{20,21} which is particularly important for working fluids.^{19,22,23} Multi-objective optimization can address multiple design criteria simultaneously, allowing for the discovery of molecules that are most fit for a specific application.

When many objectives must be optimized simultaneously, a common approach is to aggregate the objectives into a single objective function, which requires quantifying the relative importance of each objective. This method, also known as *scalarization*, reduces a multi-objective molecular optimization (MMO)



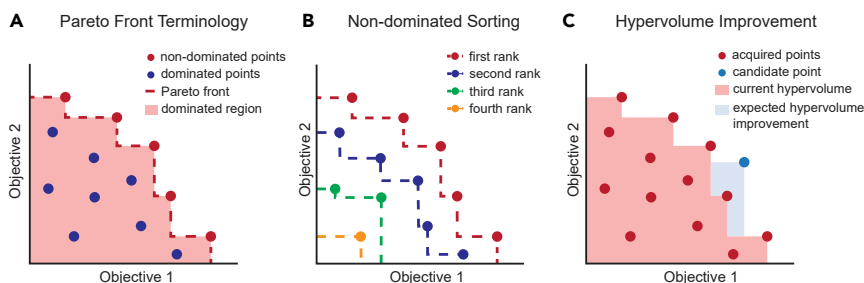


Figure 1. Terminology and acquisition functions in Pareto optimization

(A) Visual depiction of common Pareto terminology including the Pareto front, dominated and non-dominated points, and dominated region. The area of the dominated region is the hypervolume. (B) Non-dominated sorting, also referred to as Pareto ranking. (C) Hypervolume improvement for one candidate point over the current hypervolume defined by the set of previously acquired points; the drawing omits uncertainty for clarity.

problem into one that is solvable with single-objective algorithms, but the ability to explore trade-offs between objectives is limited. Furthermore, the optimization procedure must be repeated each time the scalarization function is adjusted. In contrast, Pareto optimization, which discovers a set of solutions that reveal the trade-offs between objectives, relies on no prior measure of the importance of competing objectives. This approach allows an expert to modify the relative importance of objectives without sacrificing optimization performance or repeating the optimization procedure. The solution set of a Pareto optimization contains the solution to every scalarization problem with any choice of weighting factors. For these reasons, we believe that Pareto optimization is the most robust approach to multi-objective molecular discovery.

The discovery of optimal molecules can be framed as either a search for molecules from an enumerated library or generation of novel molecules (i.e., *de novo* design).^{24,25} The extension of both discovery approaches from single-objective to multi-objective optimization has been reviewed for molecular discovery^{26,27} and more specifically drug discovery.^{28,29} However, recent developments, specifically in *de novo* design using deep learning, warrant further discussion and organization of new methods.

In this review, we organize established and emerging MMO techniques. After defining MMO and introducing relevant mathematical concepts, we describe key design choices during the formulation of an optimization scheme. Then, we provide a thorough discussion of relevant methods and case studies, first in library-based optimization and then in *de novo* design. Finally, we share some open challenges in MMO and propose future work that we believe would most advance the field.

DEFINING MMO

The molecular discovery literature is riddled with approaches to solve the inverse problem of property \rightarrow structure, many of which are labeled “multi-objective.” However, the line between MMO and single objective or constrained optimization is quite blurred. To organize the field’s communication of MMO methodologies, we classify MMO as follows:

1. Multiple objectives, which are not aggregated into a single scalar objective, are considered. Some trade-off exists between objectives (i.e., they are not perfectly correlated).
2. The domain over which to optimize (“design space”) is a chemical space. Molecules in this space may be defined either implicitly (e.g., as latent variables that can be

decoded using generative models) or explicitly (i.e., as a molecular library).

3. The goal of the optimization task is to identify molecules that maximize or minimize some molecular properties. We consider tasks that aim to identify molecules with properties within some specified range to be constrained generation, not multi-objective optimization.

Any definitive scope of MMO is bound to be somewhat subjective. Yet, we believe the preceding definition captures all relevant implementations of MMO and excludes methods that are better categorized elsewhere (e.g., as a single objective optimization or constrained optimization).

In contrast to *de novo* molecular design where molecules are proposed with few constraints other than structural validity, *pool-based* or *library-based* optimization constrains the design space to a predefined set of molecules. One approach to pool-based optimization is to predict the properties of every molecule in the set using quantitative structure-property relationships (QSPRs) and identify those that optimize the desired properties. This process is typically denoted virtual screening and can be viewed as an inefficient or brute-force approach to MMO. Virtual screening using structure-based drug design techniques for property estimation has been used to identify multi-target inhibitors^{30–32} as well as selective inhibitors.³³ In the interest of summarizing efficient optimization algorithms, we do not discuss enumeration and exhaustive screening approaches in this review.

PRELIMINARY MATHEMATICAL CONCEPTS IN MMO

The Pareto front

In MMO problems, two or more desirable molecular properties compete with one another. For *Pareto optimal* solutions, an improvement in one objective is detrimental to at least one other objective. For instance, when a selective drug is designed, strong affinity to the target and weak affinity to off-targets are both desired. However, when the binding affinities to on- and off-targets are highly correlated (i.e., they bind strongly to similar molecules), an increase in potency to the target often necessitates a decrease in selectivity. The *Pareto front* quantifies (and, in the 2- or 3-objective case, visualizes) these types of trade-offs. Figure 1A illustrates a Pareto front for two objectives that are to be maximized, with points in red representing the *non-dominated points*, which form the Pareto front and define the set of optimal solutions for the multi-objective optimization problem. For these points, an improvement in one objective

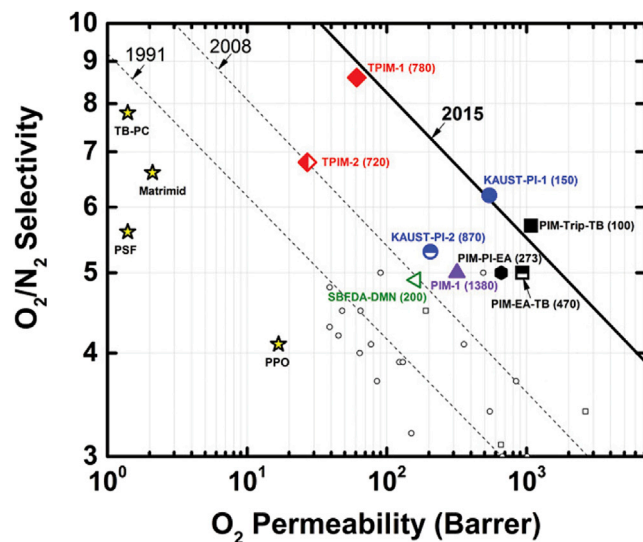


Figure 2. Progress in membranes for gas separation as revealed by the movement of a Pareto front

Reprinted with permission from Swaidan et al., *ACS Macro Lett.* 2015, 4, 9, 946–951.³⁴ Copyright 2015 American Chemical Society.

necessitates a detriment to at least one other objective. In Figure 1A, one can imagine that each objective is a desired property and that each point on the plot represents one molecule. If the two maximized objectives were measures of selectivity and binding affinity, non-dominated points represent molecules for which an improvement in binding affinity necessitates a decrease in selectivity, or vice versa. For simplicity and ease of visualization, we always consider that objectives are maximized for the remainder of the review. Pareto fronts for minimized objectives would instead appear in the lower left corner, as opposed to the upper right.

The *hypervolume* of a set is the volume spanned by the Pareto front with respect to a reference point. In the two-dimensional case, the hypervolume is the area that is dominated by the Pareto front (the red shaded region in Figures 1A and 1C). This metric can evaluate how “good” a Pareto front is: a larger hypervolume indicates a larger dominated region (i.e., a “better” Pareto front).

Progress in new materials development is often reported and visualized by the advancement of a Pareto front. As an example, in gas separation applications, membrane selectivity and permeability are two competing objectives which are both to be maximized. The trade-offs for this optimization can be visualized as a Pareto front. Figure 2 shows the improving upper bound for the two maximized objectives, which can be understood as an expansion of the Pareto front from 1991 to 2015.³⁴

Single-objective Bayesian optimization

Bayesian optimization (BO) is a strategy for black box optimization where the scalar function to be optimized, sometimes referred to as the *oracle*, may be non-differentiable or difficult to measure (costly).³⁵ The workflow of BO applied to single-objective molecular discovery is summarized in Figure 3A.

BO is an iterative optimization procedure that begins by defining some prior model to map the design space to the

objective. This model is called a *surrogate model* and, in the molecular setting, is equivalent to a QSPR model. The surrogate model is used to predict the objective values of hypothetical candidates in the design space, which an *acquisition function* uses (along with the surrogate model uncertainty) to prioritize which candidates to sample next. The newly sampled, or *acquired*, molecules are then evaluated, or *scored*, against the oracle, and these new data are used to refine the surrogate model. The process is repeated until some stopping criterion is met: the objective value of the acquired molecules converges, resources are expended, or some objective value threshold is attained.

The acquisition function is central to BO. This function quantifies the “utility” of performing a given experiment and can be broadly understood to balance both the *exploitation* and *exploration* of the design space.³⁶

In molecular BO, exploration prevents stagnation in local optima and can encourage acquisition of more diverse molecules. However, the acquisition function must also exploit by selecting candidates predicted to optimize the objective, enabling the algorithm to converge upon an optimum and identify the best-performing molecule(s). A few acquisition functions for the case where a single objective (f) is maximized are worth mentioning:

1. Expected improvement (EI):

$$EI(x) = \mathbb{E}[\max\{0, f(x) - f^*\}], \quad (\text{Equation 1})$$

in which $f(x)$ represents the objective value for some molecule x , \mathbb{E} is the expectation operator, and f^* is the best objective value attained so far from the acquired molecules.^{35,36}

2. Probability of improvement (PI):

$$PI(x) = \mathbb{E}[(f(x) - f^*) > 0]. \quad (\text{Equation 2})$$

The PI metric estimates how likely a new molecule x is to outperform the current best molecule.³⁶

3. Greedy acquisition (G):

$$G(x) = \hat{f}(x). \quad (\text{Equation 3})$$

Here, the acquisition function is simply the predicted value for the maximized objective function, regardless of uncertainty and what has been observed so far.³⁷

4. Upper confidence bound (UCB):

$$UCB(x) = \hat{f}(x) + \beta\sigma(x), \quad (\text{Equation 4})$$

in which σ is the surrogate model prediction uncertainty and β is a hyperparameter.³⁶

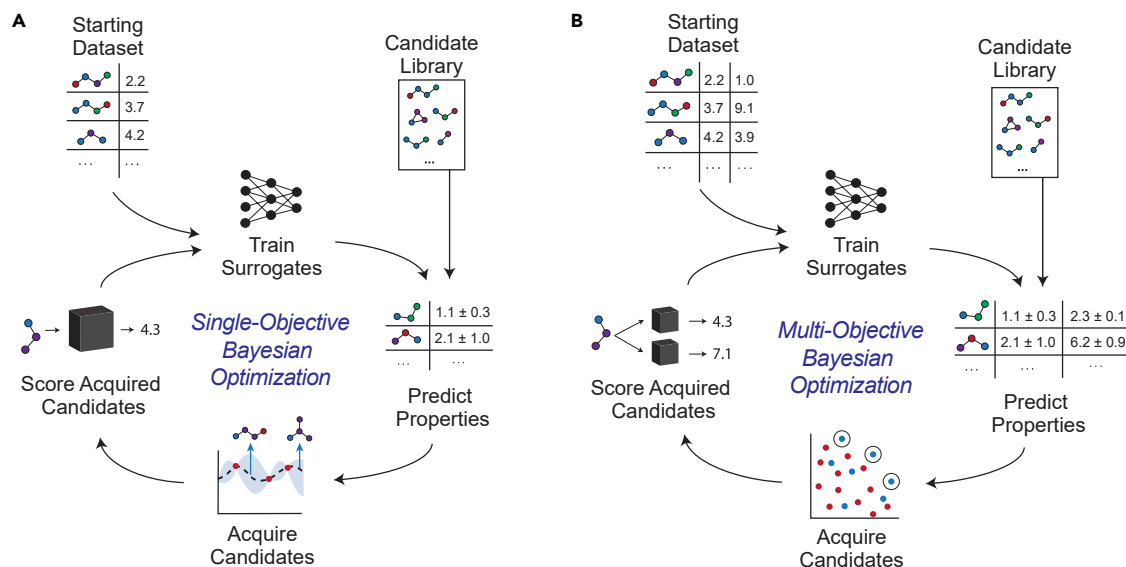


Figure 3. Overview of the Bayesian optimization workflow

There are several commonalities between the (A) single-objective and (B) multi-objective settings.

While the BO literature thoroughly discusses and tests many acquisition functions, we have only described a few that are most relevant to our discussion of MMO. We refer to readers interested in single-objective acquisition functions to Frazier's tutorial³⁵ or Shahriari et al.'s review.³⁶

Relating to BO, *active learning* can be understood as an implementation of iterative experimental design in which the objective is to train a surrogate machine learning model with greatest accuracy and fewest sampled points.³⁶ The motivation for active learning lies in the belief that a machine learning model will attain greater accuracy with less training data if it “chooses” its training data. While the iterative loop is similar to that of BO, active learning acquisition functions prioritize exploration and select points with greatest model uncertainty.

Multi-objective BO

Pareto optimization problems, in which multiple objectives are considered simultaneously without quantification of relative objective importance, must be handled with a slightly modified set of tools, although the core BO ideology remains the same (Figure 3B). First, all oracle functions must be approximated either with multiple surrogate models, a multi-task surrogate model,³⁶ or some combination thereof. Second, the acquisition function must account for all objectives without explicitly assigning a relative importance weight to each of them. Here, the goal is to expand the Pareto front, or increase the dominated hypervolume, as much as possible. We focus on three multi-objective acquisition functions:

1. Expected hypervolume improvement (EHI):

$$\text{EHI}(x) = \mathbb{E}[\max(0, \text{HV}(\mathcal{X}_{\text{acq}} \cup \{x\}) - \text{HV}(\mathcal{X}_{\text{acq}}))], \quad (\text{Equation 5})$$

in which HV is the hypervolume and \mathcal{X}_{acq} is the set of previously acquired candidates. EHI is best understood as an analog to the single-objective expected improvement, which measures improvement in hypervolume instead of objective value.

2. Probability of hypervolume improvement (PHI):

$$\text{PHI}(x) = \mathbb{E}[(\text{HV}(\mathcal{X}_{\text{acq}} \cup \{x\}) - \text{HV}(\mathcal{X}_{\text{acq}})) > 0]. \quad (\text{Equation 6})$$

PHI, comparable with probability of improvement, is the probability that an acquired point will improve the hypervolume by any amount.

3. Non-dominated sorting (NDS): NDS assigns an integer rank to each molecule by sorting the set of molecules into separate fronts. One can imagine identifying a Pareto front from a finite set of molecules (denoted first rank), removing that Pareto front, and subsequently identifying the next Pareto front (denoted second rank), as shown in Figure 1B. NDS does not consider uncertainty, and a candidate's assigned Pareto rank is taken to be its acquisition score. The first rank candidates are equivalent to the set of points that would be acquired from using greedy acquisition with every set of possible scalarization weights, so NDS can be thought of as a multi-objective analog of greedy acquisition.

Batching and batch diversity

While the canonical BO procedure evaluates candidates sequentially by acquiring the single candidate with the highest acquisition score at each iteration, many molecular oracles can be evaluated in batches. Experiments performed in well plates are naturally run in parallel, and expensive computations are

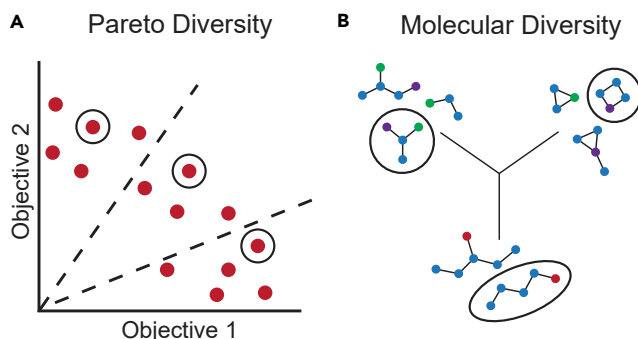


Figure 4. Contrasting definitions of diversity

During batch acquisition, one can consider one or both of (A) Pareto diversity and (B) molecular/structural diversity. Promoting one form of diversity does not necessarily improve the other.

often distributed in batches to make the best use of computational resources. In the BO workflow, this means that an acquisition function should be used to select a set of molecules, instead of just one. A naive approach, *top-k* batching, scores molecules normally and acquires the *k* candidates with the highest acquisition scores. The utility of the entire set is thus implicitly taken to be the sum of individual acquisition scores. However, the information gained from acquiring one molecule that is highly similar to another molecule in the batch is likely to be small.

In batched multi-objective optimization, the acquisition function should maximize the utility of scoring the *entire batch*. For the case of acquisition with EHI, this refers to the improvement in hypervolume after *all* molecules in a batch are acquired. One can imagine that acquiring a set of candidates very near each other on the Pareto front would not maximize this utility. An ideal batching algorithm would consider all possible batches, predict the utility of each, and select the batch with greatest utility. However, solving this combinatorial optimization exactly is intractable. Instead, approximations are used to construct batches iteratively: identify the most promising molecule, assume it has been observed, select the next most promising molecule, and repeat this until the desired batch size is achieved.³⁹

Batched optimization is more often approached with heuristics that promote some measure of *diversity* within a batch while selecting molecules with high acquisition scores. For example, the objective space can be split into regions (Figure 4A) with a limit on the number of candidates acquired in each region^{40,41}; likewise, candidates in less crowded regions along the Pareto front can be more strongly favored.⁴² Such approaches to promote *Pareto diversity* have been incorporated into multi-objective molecular design.^{43–45}

Diversity of the *design space* can also be considered during acquisition, which is distinct from Pareto diversity and can also be applied to single-objective optimization.⁴⁶ In MMO, design space diversity is equivalent to the *structural*, or *molecular*, diversity of a batch. Acquisition with structural diversity constraints might promote selection of a molecular batch with a wide array of scaffolds or functional groups (Figure 4B). Molecular diversity can be measured with metrics like Tanimoto similarity using fingerprint representations, which characterize a specific kind of structural similarity. As with Pareto diversity, structural diversity constraints can be imposed during batch acquisition in mo-

lecular optimization^{47,48} or in active learning.^{49–52} While one might predict that Pareto front diversity also indicates molecular diversity, this is not necessarily true. It is possible for two structurally similar molecules to have different properties and therefore lie in different regions of the objective space; conversely, molecules with similar properties are not necessarily structurally similar.

FORMULATING MOLECULAR OPTIMIZATION PROBLEMS

A molecular optimization task always begins with some statement of desired properties. Some of the subsequent formulation decisions are listed in Figure 5. First, the individual properties must be converted to mathematical objectives. Then, the means of proposing candidate molecules, either *de novo* or library based, must be selected. If more than one objective exists, they must either be aggregated into a single objective or treated with an appropriate multi-objective formulation. Finally, an acquisition function, or selection criterion in the case of *de novo* design, must be selected. In this section, we explore some of these design choices in detail.

Converting a desired property to a mathematical objective function

In the formulation of any MMO task, after properties of interest are identified by a subject matter expert, the individual objectives must be quantitatively defined (Figure 5, panel 2). While this seems like an easy task, framing the objectives can be subjective in nature. If one property of interest for a molecular optimization task is estimated by a score $S(x)$, there are still multiple ways to represent the corresponding value to be maximized ($J(x)$), including but not limited to:

1. A continuous, strictly monotonic treatment, where a greater value is strictly better:

$$J(x) = S(x). \quad (\text{Equation 7})$$

2. A thresholded, monotonic treatment, where some minimum T is required:

$$J(x) = \begin{cases} S(x) & S(x) \geq T \\ -\infty & S(x) \leq T \end{cases} \quad (\text{Equation 8})$$

3. A Boolean treatment, where some minimum is required and no preference is given to even higher values:

$$J(x) = \begin{cases} 1 & S(x) \geq T \\ 0 & S(x) \leq T \end{cases}. \quad (\text{Equation 9})$$

The most appropriate representation depends on the property of interest and the application, demonstrated here for common properties of interest for novel drug molecules. If S predicts a

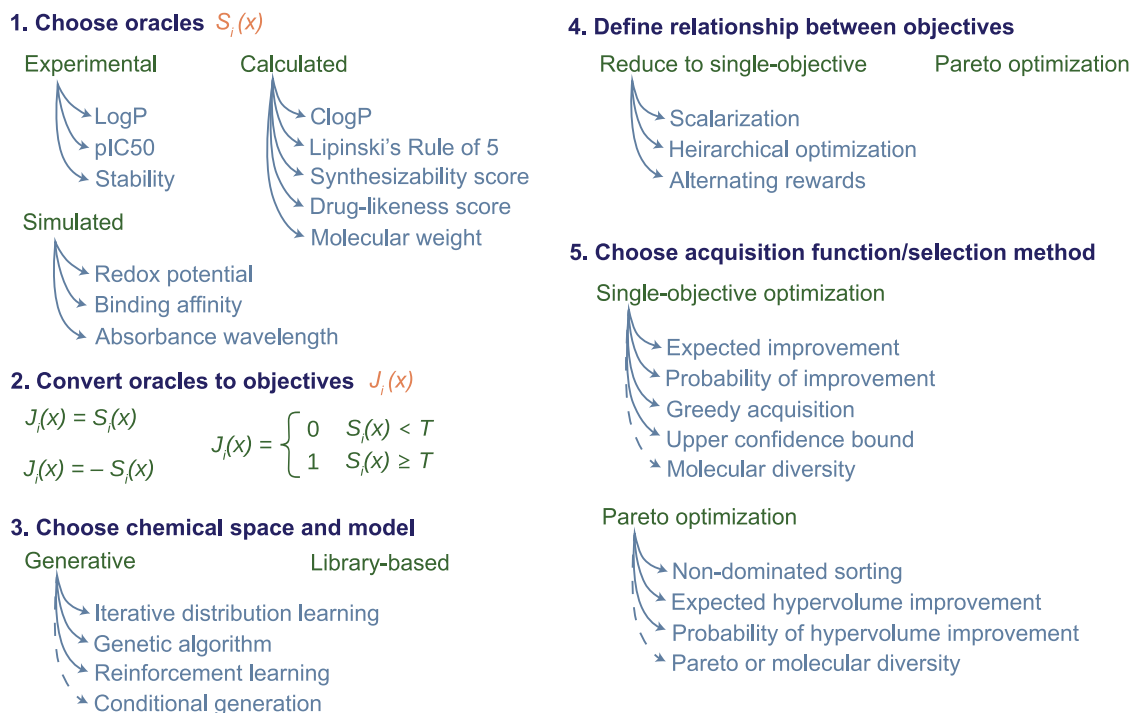


Figure 5. Decisions when formulating MMO problems

Iterative generative models tend to employ selection criteria for retraining or propagation, which are analogous to acquisition functions in Bayesian optimization. As discussed in later sections, conditional generation aims to propose molecules with a specified property profile in a non-iterative manner and therefore does not utilize selection criteria or an acquisition function. Single-objective acquisition functions can only consider molecular diversity, while Pareto acquisition functions can consider both molecular and Pareto diversity.

ligand's binding affinity to a target protein, a higher affinity is often better, so the first representation may be most appropriate. If S predicts solubility, there may be no additional benefit of greater solubility once a certain solubility is met that allows for sufficient delivery and bioavailability. In this case, the third representation, which is most consistent with a property constraint instead of an optimized objective, would be most fitting. In a similar manner, remaining components of Lipinski's Rule of 5⁵³ define some threshold, and no extra benefit is attained once the threshold is met. These heuristics may be most appropriately defined as constraints and not optimized objectives.

The perspectives of domain experts during objective formulation are extremely valuable to ensure that molecules identified as optimal are suitable for the application. However, in cases where expertise is not available or a specific threshold is unknown, we argue that solving the problem with a simple continuous representation (representation 1) is most robust because it requires no predefined hyperparameters or assumptions. This way, constraints can later be imposed on the solution set without needing to repeat the optimization from scratch.

Choosing between library-based selection and *de novo* design

Once the objectives are defined, an approach to chemical space exploration must be chosen. The scope of exploration can be limited to an explicitly defined molecular library, which can be constructed to bias exploration toward chemical spaces relevant to a specific task. Alternatively, a *de novo* design tool can be

used to ideate novel molecules not previously seen or enumerated. The type of generative model influences the area of chemical space that is explored.⁵⁴ For example, the chemical space explored by genetic algorithms (GAs) may be constrained by the molecules used as the initial population and the set of evolutionary operators that are applied to the population. In a more general sense, the molecules that can be generated by any *de novo* model will be determined by the training set and many other design choices. Broadly, constraining the design space to a carefully enumerated chemical library can bias exploration toward molecules that are synthesizable and/or more fit for a given application, potentially sacrificing the creativity and theoretically wider chemical space accessible with *de novo* design.⁵⁴

Defining the relationship between different objectives

Once individual objective functions are defined and the chemical space approach is chosen, the next challenge is to decide how to consider all objectives simultaneously. The most naive choice is to simply combine the objective functions into one aggregated objective function, referred to as *scalarization*. The scalarized objective function is most commonly a weighted sum of objectives,^{55–61} with weighting factors indicating the relative importance of different objectives. A weighted sum of multiple binding affinities has been used to identify multi-target⁶² as well as selective inhibitors.⁶³ Nonlinear scalarization approaches are also utilized in MMO problems.^{64–66} For example, Cardoso Gajo et al. divide predicted drug activity by toxicity to yield a scalarized objective function.⁶⁷ The objective function can also be framed

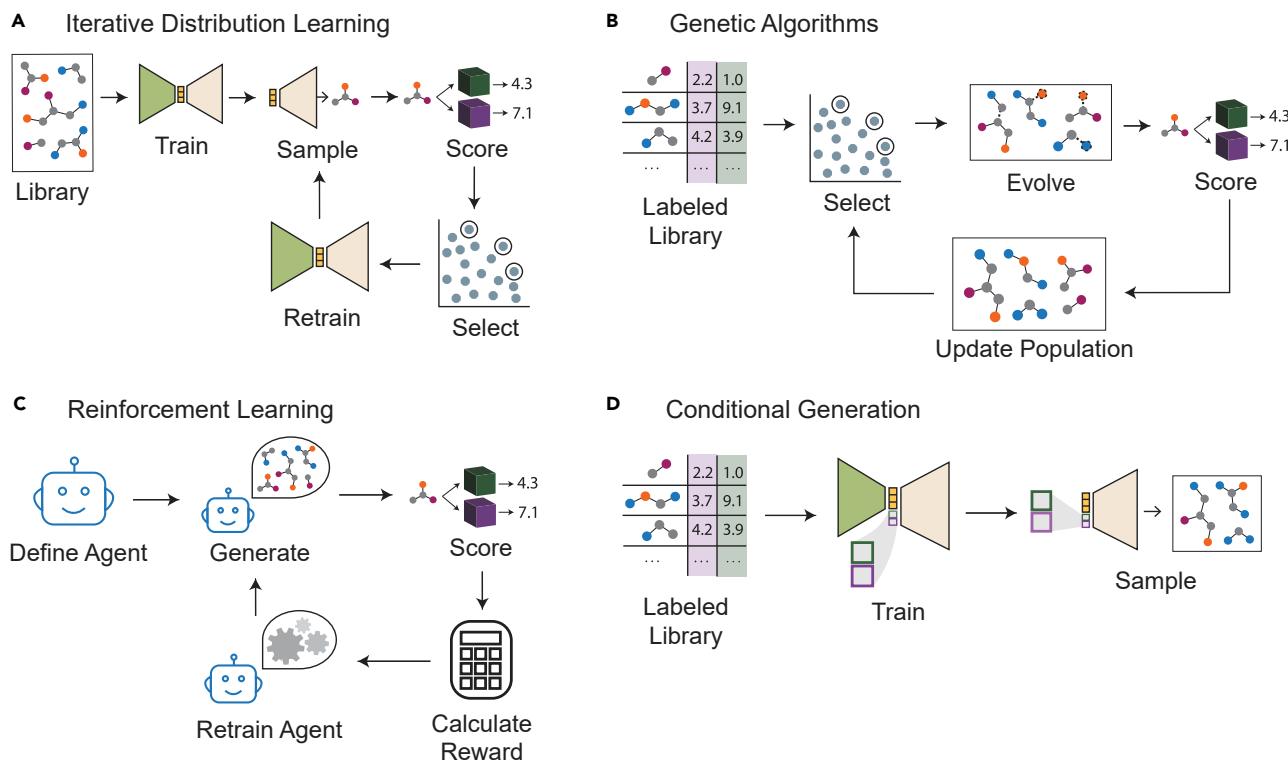


Figure 6. Optimization workflows for various generative model categories

Note that all model classes, except conditional generation, involve a scoring step and are designed to be iterative. The reward calculation step in reinforcement learning and the selection step in distribution learning and genetic algorithms are analogous to an acquisition function in multi-objective Bayesian optimization. While the termination criterion is not explicitly shown for distribution learning, genetic algorithms, and reinforcement learning, these iterative loops can accommodate various stopping criteria. We also emphasize that while an autoencoder architecture is depicted in both distribution learning and conditional generation, these generators can also be recurrent neural networks or other generative architectures.

as a product of Booleans,⁶⁸ each of which denotes whether a given threshold is met. This scalarization approach has been utilized to identify multi-target kinase inhibitors.⁶⁹

Booleans can also be summed to define an objective function, commonly referred to as multi-property optimization.⁷⁰ Desirability functions, another common scalarization method, transform one or more objectives into a scale-free value between 0 and 1.^{71–73} As with the definition of individual objectives, the scalarization function must be justified by the use case. There are alternatives to scalarization that also reduce a multi-objective optimization into one that can be solved with single-objective algorithms, such as defining a hierarchy of objective importance⁷⁴ or using alternating rewards to maximize each objective in turn.^{75,76}

However, the solution to a scalarized multi-objective problem is equivalent to just a single point out of the many non-dominated solutions that exist on the Pareto front. Scalarization is overly simplistic and requires a user to quantify the relative importance of different objectives. It therefore fails to inform a user about the trade-offs between objectives. Even when the relative importance of objectives is known or can be approximated *a priori*, scalarization is strictly less informative than Pareto optimization, which identifies the full set of molecules that form a Pareto front. We focus exclusively on Pareto optimization approaches to molecular discovery throughout the remainder of this review.

EXAMPLES OF MMO FROM VIRTUAL LIBRARIES

Library-based MMO aims to identify the Pareto front (or a set close to the Pareto front) of a large molecular library while scoring few molecules with the objectives. The well-established BO workflow (Figure 3B) is exemplified by the retrospective studies of del Rosario et al.⁷⁷ and Gopakumar et al.⁷⁸ In general, the iterative optimization scheme entails training a surrogate model to predict properties of interest, selecting molecules for acquisition using surrogate model predictions and uncertainties, scoring the acquired molecules with the ground-truth objectives, and retraining the surrogate model.

Janet et al.⁴⁷ apply this methodology to discover transition metal complexes for redox flow battery applications with maximized solubility and redox potential. Ideal complexes must be soluble in polar organic solvents commonly used for flow batteries and have high redox potentials to yield sufficient cell voltage. The design space the authors explore is a combinatorial library of almost 3 million complexes. A neural network surrogate model predicts solubilities and redox potentials from feature vector representations of complexes.⁷⁹ DFT calculations served as the oracle for both solubility and redox potential, and the expected hypervolume improvement acquisition function was used. To encourage exploration of structurally diverse complexes, the top 10,000 performers according to EHI were clustered in feature

space to identify and evaluate 100 medoids. Improvements of over three standard deviations from the initial random set of complexes were observed for both objectives in just five iterations, which the authors estimate to represent a 500X reduction in simulations compared with a random search.

In a similar vein, Agarwal et al.⁴⁴ use library-based Pareto optimization to search for redox-active materials with minimized reduction potential and solvation free energy. A third objective penalized deviation from a target peak absorption wavelength of 375 nm. Candidates were scored with expected hypervolume improvement, while crowding distance constraints ensured acquisition of a diverse set along the Pareto front. When retrospectively applied to a dataset of 1,400 molecules, a random search required 15X more evaluations than did BO to acquire molecules dominating 99% of the total possible hypervolume. Then, a prospective search was performed on a set of 1 million molecules, with the prior dataset serving as the first set of acquired molecules. Of the 100 molecules acquired during prospective BO iterations, 16 new Pareto-optimal molecules were identified.

Most pool-based MMO problems follow this exact workflow with minor variability in the choice of acquisition function and consideration of diversity. This approach works effectively and is almost guaranteed to outperform random search baselines. While there is certainly room for algorithmic improvement (e.g., increasing sample efficiency of surrogate models, exploring the effects of batch size and diversity), we expect that future work will largely focus on additional applications incorporating more meaningful objective functions and experimental validation.

EXAMPLES OF MMO USING GENERATIVE MODELS

The primary drawback of pool-based MMO is the explicit constraint on the chemical space that can be accessed. *De novo* design relaxes this constraint and can, in principle, explore a wider region of chemical space. At the same time, *de novo* design can introduce additional challenges related to chemical validity⁸⁰ and synthesizability.⁸¹ In many generative models, molecules are proposed as SMILES/SELFIES strings, graphs, or synthetic pathways. Some generate novel molecules by decoding continuous embeddings into discrete molecular structures while others modify those already identified with discrete actions. We focus not on the details of each model, but instead on how certain categories of models aid in the molecular optimization task. A reader interested in a detailed discussion of generative models, which is outside the scope of this review, is directed to other publications.^{80,82–84}

The myriad of multi-objective *de novo* design approaches noticeably lack standardization. Unlike library-based discovery where multi-objective optimization is a modest extension of BO, the adaptation of generative models to MMO is not nearly as straightforward. We therefore introduce another categorization scheme for case studies in this section, as summarized in Figure 6.

Iterative retraining for distribution learning

Generative models that are designed for *distribution learning* are intended to ideate molecules exhibiting a distribution of struc-

tures similar to those of the training set.⁸⁵ A very basic approach to optimization with an unsupervised generative model is to sample a set of molecules, evaluate their properties, and identify those that optimize the objective function; to extend this to multi-objective optimization, the Pareto front of the sampled set can be identified by evaluating all oracles.⁸⁶ This approach essentially uses a generative model to define a virtual library suitable for exhaustive screening. Optimization schemes can use distribution learning *iteratively* to progressively shift the distribution of generated molecules and push the Pareto front. To achieve this, generative models are iteratively retrained on the increasingly promising (e.g., closest to the Pareto front) subsets of the molecules they propose. This process is akin to a simulated design-make-test loop, in which *design* is analogous to sampling, *make* to decoding to a molecule, and *test* to evaluating the oracles.

The iterative distribution learning workflow for single-objective optimization is exemplified by the library generation strategy defined by Segler et al.⁸⁷ to identify inhibitors predicted to be active against the 5-HT_{2A} receptor. Here, a subset of molecules from the ChEMBL database, with corresponding experimental pIC₅₀ values against 5-HT_{2A}, was used to train both a SMILES-based recurrent neural network (RNN) and a QSAR classifier to predict whether a molecule inhibits 5-HT_{2A}. Then, sequences of characters were randomly sampled from the RNN to generate SMILES representations of novel molecules. Molecules predicted by the QSAR classifier to be active were used to retrain the model, progressively biasing the generator to propose active molecules. After four iterations of retraining, 50% of sampled molecules were predicted to be active, a significant increase from only 2% in the initial random library. The same procedure has also been employed using a variational autoencoder to generate molecules with high docking scores to the DRD3 receptor.⁸⁸

The extension of the method to multiple objectives is best illustrated by Yasonik⁸⁹ for the generation of drug-like molecules. As before, a recurrent neural network was pretrained to generate valid molecular SMILES strings. Five oracles associated with drug-likeness were then minimized: ClogP (estimated lipophilicity), molecular weight, number of hydrogen bond acceptors, number of hydrogen bond donors, and number of rotatable bonds. A set of about 10k novel, unique, and valid molecules were sampled and scored according to the five properties. NDS was used to select half of these molecules for retraining. The use of NDS distinguishes this Pareto optimization from Segler et al.'s single-objective optimization. Although continuous objective values were used during selection of molecules for retraining, constraints associated with the oracles, derived from the "Rule of Three"⁹⁰ (an extension of Lipinski's Rule of 5⁵³), were used to evaluate the generator's performance. After five retraining iterations, the fraction of molecules that fulfilled all five constraints increased from 2% to 33%. While there is no evidence that the Pareto front was shifted outwards (i.e., that the dominated hypervolume increased) after retraining iterations, this study demonstrates that a generative model's property distributions for multiple objectives can be shifted simultaneously.

In addition to recurrent neural networks, as in the prior two examples, variational autoencoders and other generative models can be iteratively retrained to simultaneously fulfill

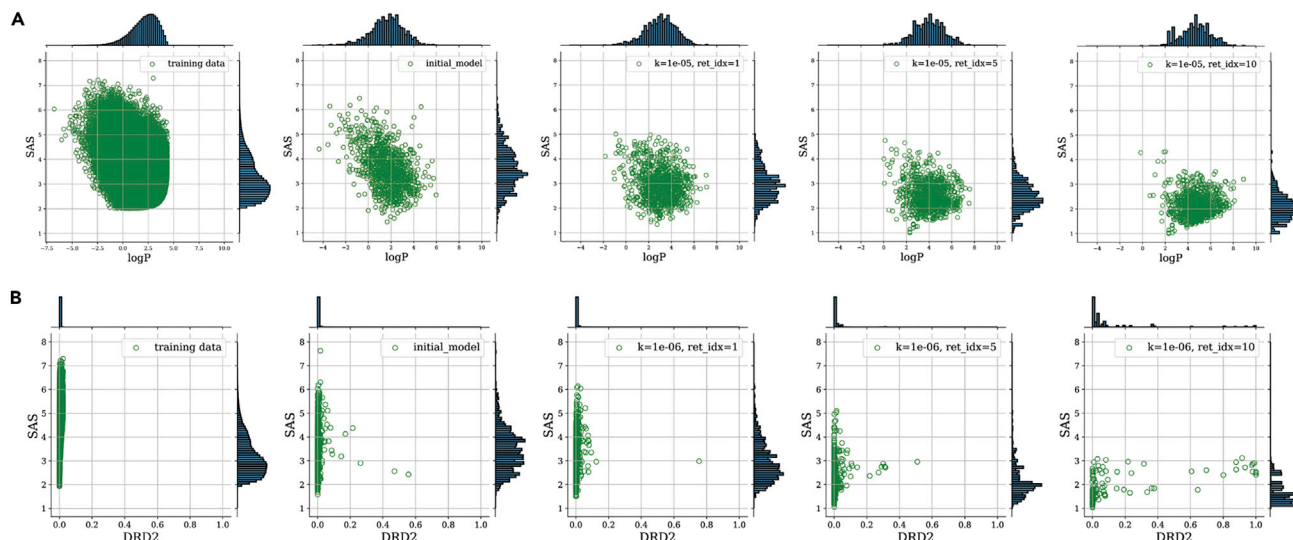


Figure 7. Advancement of the Pareto front from Abeer et al. using iterative retraining for distribution learning

Both (A) and (B) are from the same optimization task, with each set only showing two objectives for ease of visualization. The first and second columns are the distribution of the training molecules and the first batch of sampled molecules, respectively. The following 3 columns depict molecules sampled from the model after 1, 5, and 10 iterations. Reproduced from Abeer et al.⁹² with permission from the authors.

multiple property constraints.⁹¹ Abeer et al.⁹² describe one such approach to generate drugs with high predicted binding affinity to the DRD2 receptor, high ClogP, and low synthesizability score using a VAE as the unsupervised generator. After initial training, sampling, and scoring, the best molecules were selected according to their Pareto rank, but some random molecules were also included in the retraining set. Importantly, the authors show a progression of the two-dimensional Pareto fronts beyond those of the original training set: they identified molecules that are strictly superior to (i.e., that “dominate” in a Pareto optimality sense) the best molecules in the training set. Two such plots are shown in Figure 7. Here, it is clear that this method is capable of increasing the dominated hypervolume and identifying novel molecules that have property values outside of the objective space spanned by the training set.

Genetic algorithms

In contrast to many deep learning architectures, genetic algorithms (GAs) do not rely on a mapping between continuous and discrete spaces. Instead, molecules are iteratively transformed into new ones using *evolutionary operators* like mutations and crossovers. Molecular *mutations* may include the addition or removal of atoms, bonds, or molecular fragments, while molecular *crossover* involves molecular fragment exchange between two parent molecules. GAs begin with a starting population of molecules that are scored by the oracle function(s). Selection criteria are imposed to determine which molecules in the population are chosen as *parents* to be propagated. This selection step is what guides a GA to optimized molecules and, like an acquisition function in BO, determines whether an optimization is a Pareto optimization or not. Evolutionary operators are randomly chosen and applied to the parents, and the population is updated with the resulting molecules.

GAs were the first popularized polymer⁹³ and small-molecule⁹⁴ generators. In 1995, Sheridan and Kearsley⁹⁴ pro-

posed generating small molecules by iteratively evolving integer sequence representations of molecules. That same year, Weber et al.⁹⁵ used a GA to find optimal molecules from a synthetically enumerated library. Since then, GAs have adopted evolutionary operators which function directly on molecular graphs^{96–98} or SMILES strings.⁹⁹ Some GAs even mutate molecules using chemical reaction templates to encourage synthesizability.^{95,100,101} Multiple objectives can be scalarized during selection to frame a multi-objective GA as a single objective one.^{96,98,102,103}

As with any generative model, if the selection criteria consider multiple objectives simultaneously without imposing assumptions about relative importance, a GA can advance the population’s Pareto front. One such GA was proposed by Brown et al.⁹⁷ to generate “median molecules,” which maximize Tanimoto similarity¹⁰⁴ to two different molecules simultaneously. In each iteration, molecules in a population are manipulated with either mutations (add/delete atoms, add/delete bonds) or crossovers (molecular fragment exchange between two parent molecules). NDS, using the two Tanimoto similarities as objectives, determine which molecules are selected for propagation. The critical adaptation for the multi-objective case is the use of Pareto ranking—specifically, NDS—as a selection criterion, instead of using a single property estimate or a scalarization of multiple properties.

A comparable multi-objective GA, presented by Nicolaou et al.,⁴⁸ generates ligands with maximized docking scores for a target receptor (estrogen receptor β) and minimized scores for a negative but closely related target (estrogen receptor α). As an extension from the prior example, the NDS selection criterion was modified to include *nicheing* and *elitism*. Nicheing encourages structurally diverse populations by grouping candidates into *niches* based on their structural similarity during selection, and only a set number of molecules may be acquired in each niche. Promoting diversity can be especially beneficial to GA

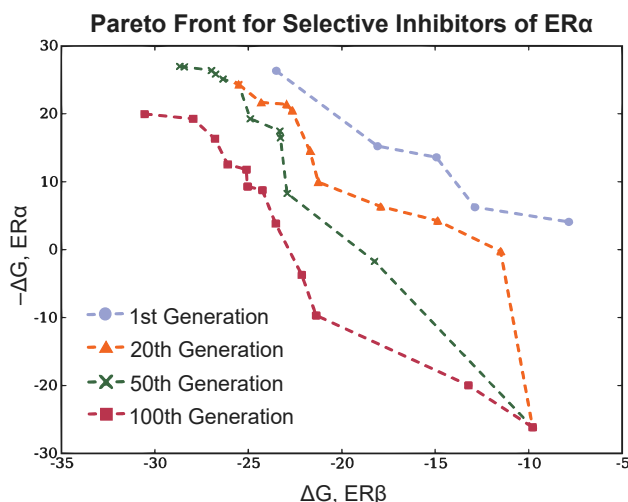


Figure 8. Pareto front for the identification of selective inhibitors

The ΔG values represent docking scores. Note that the Pareto front in this plot is located in the bottom left. The Pareto front is shown after 1, 20, 50, and 100 iterations. It clearly shifts to the bottom left with each iteration. Here, niching is used but elitism is not. Redrawn from Nicolaou et al.⁴⁸

performance, as GAs are constrained by their starting set and set of modification operators.^{105,106} When elitism is imposed, all Pareto-dominant molecules found during prior iterations are appended to the population before selection to prevent good molecules from being “forgotten.” The authors report that both elitism and niching improve optimization performance. The depicted progression of the Pareto front is replicated here (Figure 8). The notion of optimizing against a negative target can be generalized into a “selectivity score” that aggregates affinity to multiple off-target controls.¹⁰⁷ Properties related to bioavailability, such as blood-brain barrier penetration and aqueous solubility, can serve as additional objectives for the discovery of novel drug molecules, as reviewed by Ekins et al.²⁸

The effect of diversity-aware acquisition is further explored by Verhellen,⁴³ wherein the effectiveness of two different multi-objective GAs that promote Pareto diversity are compared. Both GAs use NDS to select the population members to be propagated as parents of the next generation. The first, NSGA-II,⁴² promotes selection of molecules with a larger distance from other molecules in the objective space and has precedent in application to a synthesizability-constrained molecular GA.¹⁰¹ The second, NSGA-III,⁴¹ enforces diversity by requiring at least one molecule to be acquired in each of a set of reference regions in the objective space (Figure 4A). Both GAs are applied to seven molecular case studies, each with a different set of objectives including affinity to a target, selectivity, and/or molecular weight. Using the dominated hypervolume as an evaluation metric, both multi-objective optimization approaches outperform a weighted-sum scalarization baseline, but there is no clear winner among the two NSGA algorithms. A measure of internal similarity indicates that the structural diversity decreased with each evolutionary iteration. Nonetheless, the selection criteria promoted Pareto diversity, demonstrating that Pareto diversity can be achieved without necessarily requiring molecular, or structural, diversity.

Reinforcement learning

Reinforcement learning (RL)-based generative models are trained to create molecules by learning to maximize a reward function quantifying the desirability of generated molecules. In molecular RL, a *policy* determines which molecules are generated and can be iteratively updated to maximize the reward as new molecules are generated and scored. The set of actions or choices available to the policy is denoted the *action space*. The framing of the reward function, analogous to the BO acquisition function and GA selection criteria, determines whether an RL method utilizes Pareto optimization.

When the learned policy generates molecules by modifying a previous population of molecules, the action space may be comprised of atom- and bond-level graph modifications^{108–110} or a set of fragment-level graph modifications.¹¹¹ In a similar manner, graph modifications resulting from chemical reactions can constitute the action space to promote synthesizability.¹¹² When the policy is a deep learning generator that designs molecules from scratch, any *de novo* generator that decodes latent variables to a molecule, such as SMILES recurrent neural networks, can be considered as the policy.^{113–117} Typically, these policies are trained using policy gradient algorithms (e.g., REINFORCE).¹¹⁸

Most RL approaches to molecular discovery, and specifically to drug design,^{119,120} optimize a reward that considers a single property^{113–115} or a scalarized objective.^{30,59,92,108–112,116,121–126} We are aware of only one molecular RL approach whose reward function directly encourages molecules to be generated along a Pareto front. In DrugEx v2, presented by Liu et al.,¹²⁷ RL is used to generate multi-target drug molecules. To promote the discovery of molecules along the Pareto front, NDS is used to calculate the reward. The authors test their algorithm with both this Pareto reward function and a weighted sum reward function. In the weighted-sum benchmark, the weighting factors were set as dynamic parameters that were altered during inference to encourage the model to find solutions at different locations on the Pareto front, analogous to the alternating reward approach to scalarization. For the multi-target discovery case, the fraction of generated molecules with all properties above some threshold values was 81% with the Pareto scheme and 97% with the weighted sum scheme. The two approaches were only compared in this constraint-style evaluation, not in terms of a Pareto optimization criterion such as hypervolume improvement, so it is not clear if the lackluster performance of the Pareto optimizer is merely due to this misalignment of evaluation criteria.

Conditional generation

Conditional generators produce molecules that are meant to achieve some set of user-defined properties instead of directly maximizing or minimizing them in an iterative manner. These models generate one or more molecules predicted to fulfill the specified properties in one pass and without any iterative feedback as in previously discussed approaches. Although our focus in this review is on multi-objective optimization, we feel that discussing the role of conditional generators in MMO is necessary due to their prevalence in the field and the ease of extending from single-objective (single-constraint) conditional generators to multi-objective (multi-constraint) conditional generators.

Many conditional generators are autoencoders that map molecules to latent embeddings and vice versa. To generate molecules with specific properties, the latent variables of these generators can be manipulated during training such that they represent the properties of interest. One such manipulation applied to variational autoencoders is to recenter the prior distribution around the associated molecule's property value c instead of the origin, encouraging the latent distribution to match $\mathcal{N}(c, \sigma^2)$ instead of $\mathcal{N}(0, \sigma^2)$.^{128–130} This approach can be expanded to multiple objectives by centering each latent dimension along a different property of interest.¹²⁸ Then, during inference, sampled latent variables are chosen according to the desired property values with at least partial success.

Autoencoders can also be manipulated for conditional generation by directly feeding the property value(s) of training molecules to the decoder during training.^{131,132} As one example, Lim et al.¹³³ use this approach to fulfill certain “drug-like” property criteria. During CVAE (conditional VAE) training, a condition vector including molecular weight, ClogP, number of hydrogen bond donors, number of hydrogen acceptors, and topological polar surface area is appended to the latent space during decoding. Then, during generation, a manually specified conditional vector influences the decoder to generate molecules with the stated properties. In all case studies, less than 1% of generated molecules have properties within 10% of the values set in the condition vector. Another study using a similar architecture¹³⁴ demonstrates that it is possible for the properties of up to 33% of generated molecules, “when rounded up,” to reflect the specified properties. In this case, it appears that this fraction strongly correlates with how many training molecules also fulfilled those constraints.

Some conditional generators modify existing molecular graphs or scaffolds provided as input instead of generating molecules from scratch. These models are typically trained with matched molecular pairs: pairs of molecules with only one well-defined structural transformation that causes a change in molecular properties.^{135,136} One such single-objective generative model is intended to “translate” molecules that are inactive as DRD2 inhibitors to active inhibitor molecules,¹³⁷ wherein activity is predicted by a trained classifier. The generative model is presumed to learn graphical translations that most contribute to inhibitory strength. This methodology can be extended to the multi-constraint case if improvements in multiple properties are desired.^{138–140} For example, MolGPT, a conditional generator proposed by Bagal et al.,¹⁴¹ accepts a scaffold and desired property values. It then outputs a molecule that it believes to fulfill the input constraints. Molecules are completed from scaffolds as SMILES strings, and the model is trained on sets of {scaffold, molecule, properties}. The success of MolGPT in meeting target properties relies on having molecules with that property be well-represented in the training set. While MolGPT is able to generate molecules conditioned on multiple properties, the authors do not report whether their model is capable of generating molecules with combinations of property values not present in the training set.

The effectiveness of conditional molecule generators depends not only on their ability to generate valid and unique molecules, but also on the accuracy of the implicit molecule-property model. If this model is inaccurate, the generator will suggest mol-

ecules that do not actually exhibit the desired properties. We further emphasize that, in order to identify Pareto-optimal molecules, the model must be able to extrapolate past the training set because, by definition, Pareto-optimal molecules have properties (or combinations of properties) that are not dominated by members of the training set.

Therefore, we find it unlikely that these non-iterative conditional generators will succeed in advancing the Pareto front. This is in contrast to iterative optimization methods, wherein the predictive capability of the generators is improved for newly explored regions of chemical space with each iteration.

Furthermore, the nature of conditional generators requires that a user know what property value ranges are feasible. Based on the discussed and other case studies,^{142,143} conditional generators perform well primarily when attempting to generate novel molecules with property combinations spanned by the training set. A pIC₅₀-conditioned model would propose some set of molecules if asked to achieve a pIC₅₀ value of 100, even though such a value is unrealistic. Their behavior in these settings is not well understood, so a user may need to know which property constraints are valid or possible. Due to these concerns, we caution the reader that conditional generators may not be most appropriate for Pareto optimization tasks.

Hybrid approaches

The case studies that we have shared so far fall neatly into our defined categories. However, certain other approaches that combine methods from multiple categories or otherwise deviate from this classification are worth mentioning.

Grantham et al.⁴⁵ introduce one such hybrid approach, in which latent representations of molecules are mutated with a GA and decoded to generate new molecules. A variational autoencoder is first trained to encode molecules into latent vectors. After encoding the starting population, mutations are applied to their corresponding latent vectors, which are then decoded. From this new set of evolved molecules, NDS with a crowding distance constraint (specifically, NSGA-II⁴²) is used to select new molecules to use for retraining the autoencoder. The proposed method outperforms two BO baselines in terms of the hypervolume of the final Pareto front when applied to an optimization of ClogP, QED, and synthesizability score. A similar methodology was used to optimize both drug-likeness properties and binding affinity (estimated via docking scores) to carbonic anhydrase IX.¹⁴⁴

Iterative retraining has also been used to improve the performance of a conditional generator. In one example, a conditional graph generator is fine-tuned with molecules that are active against both JNK3 and GSK3 β .¹⁴⁵ This workflow essentially follows the iterative retraining of distribution learning algorithms, but uses conditional generation to provide an extra bias toward sampling molecules with favorable properties. In a similar manner, RL methods can be considered conditional generation if the reward function favors molecules with a target property profile.^{146–148} Two such methods^{69,149} use RL to generate molecules that are predicted to be dual inhibitors of GSK3 β and JNK3 receptors according to pretrained surrogate models. In the final populations in both studies, 100% of molecules are active against both inhibitors. However, the dataset used in both studies for training already includes a small fraction of dual inhibitors.

Therefore, discovering “active inhibitors” in this case is equivalent to discovering the chemical space that is classified as active according to the surrogate models, and this task is easier than extrapolating with a continuous oracle. In general, the reported success of generators conditioned on Boolean values (instead of continuous ones) can be overoptimistic, as the degree of optimization success is harder to quantify with metrics such as the hypervolume.

DISCUSSION

In the description of library-based MMO, we explained that these methods are a natural extension of BO. In contrast, *de novo* methods stray farther from classic BO, although some aspects of BO acquisition functions are present in generative workflows. In particular, NDS is often used as the selection criterion for re-training (distribution learning) or propagation (GAs). Other conventional BO acquisition functions, such as EHI and PHI, are rarely incorporated into optimization with generative models. These acquisition functions use the uncertainty in surrogate model predictions, which aids in the balance between exploration and exploitation. But most generative optimization architectures score molecules with the ground truth objectives during selection, thus bypassing uncertainty quantification and making EHI and PHI unusable as acquisition functions. An opportunity exists to incorporate Bayesian principles into *de novo* design by including a separate surrogate model that predicts objective function values and can be retrained as new data are acquired to guide selection. These and other adjustments to *de novo* optimization approaches may help bridge the gap between generation and model-guided optimization.

We have also observed that the performance of Pareto optimization approaches is often evaluated using individual property values or constraints. These metrics, however, reveal little about the *combination of properties* of discovered molecules, which is of foremost interest in MMO. Hypervolume improvement can indicate the shift in the Pareto front, but other qualities of the discovered molecules related to the Pareto front^{150,151} can be of equal importance, including the density of the Pareto front or the average Pareto rank of the molecules. In molecular discovery, imperfect property models are often used as oracles. In these cases, it is beneficial to discover a dense Pareto front and many close-to-optimal molecules *according to QSPR predictions*, even if not all increase the hypervolume. Naturally, some molecules that are predicted to perform well will not validate experimentally, and having a denser population to sample from will increase the probability of finding true hits. For the same reason, promoting structural diversity and not just Pareto diversity is a way to hedge one's bets and avoid the situation where none of the Pareto-optimal molecules validates.

In batched multi-objective optimization, Pareto diversity can be considered during acquisition to promote exploration. In molecular optimization, structural diversity similarly encourages exploration of a wider region of chemical space. Thus, in MMO, both potential measurements of diversity are relevant, and either or both can be used during optimization. At this point, neither diversity metric has been shown to outperform the other in MMO tasks, and the question of how best to incorporate both

into acquisition (or whether this actually benefits optimization) remains. At present, diversity-aware acquisition is more commonly incorporated into multi-objective GAs rather than other optimization architectures like RL or iterative distribution learning, although diversity-promoting acquisition functions may improve the performance of these generative models as well.

We have argued that Pareto optimization is a more practical approach to many molecular discovery tasks than scalarization or constrained optimization, but the ability of Pareto optimization to scale to several dimensions must also be addressed. NDS increasingly fails to differentiate the optimality of solutions with more objectives, as more and more points are non-dominated in a higher-dimensional space.¹⁵² The numerical estimation of hypervolume has a computational cost that scales exponentially with the number of objectives, making EHI and PHI acquisition functions also increasingly difficult to use in high dimensions.¹⁵² The increased computational costs associated with fine-tuning many surrogate models and scoring candidates for every objective contribute to scalability issues as well. Considering the challenges faced with Pareto optimization of many (more than three) objectives, scalarizing certain objectives or converting some to constraints to make the problem solvable may be the most practical approach, especially when some objectives are known to be more important than others. The question of whether Pareto optimization can robustly scale to many objectives is a worthwhile one only if a problem cannot be feasibly reduced. The visualization of the Pareto front is an additional consideration; objective trade-offs are more easily conveyed with a Pareto front of two or three objectives. Ultimately, the optimal formulation of an MMO problem will depend on the use case, and collaboration with subject matter experts can ensure that the problem formulation is feasible but does not impose unrealistic assumptions.

Beyond these unique challenges posed by multi-objective optimization, many challenges from single-objective optimization remain relevant.^{25,80,153} The first is the need for realistic oracle functions that can be evaluated computationally but meaningfully describe experimental performance; this is closely related to the need for more challenging benchmarks to mimic practical applications. Optimizing QED, ClogP, or a Boolean output from a classifier are easy tasks and are not good indicators of robustness or generality. Generative models specifically must also prove effective with fewer oracle calls, which is often the bottleneck when molecules must be scored with experiments or high-fidelity simulations.¹⁵⁴ Active learning and optimal experimental design techniques¹⁵⁵ will continue to inspire methods that improve molecular optimization sample efficiency.^{49,52} For experimental applications, the synthesizability of generated molecules is an additional factor that must be considered⁸¹ and can be cast as a continuous objective or a rigid constraint. Experimental prospective validation is essential to demonstrate the viability of molecular discovery algorithms, although algorithmic advances can be made more rapidly with purely computational studies.

Conclusion

Although many approaches to computer-aided molecular design have been developed with just single-objective optimization in mind, molecular discovery is a multi-objective optimization problem. In certain situations, such as optimization from a library (BO-accelerated virtual screening), the extension from

single-objective to multi-objective requires only minor modifications, e.g., to the acquisition function and to the number of surrogate models. In contrast, *de novo* design workflows vary more in methodology and are less directly analogous to BO. The use of Pareto rank as a reward (for RL) or the use of NDS to select sampled molecules to include in subsequent populations (for GAs) or training sets (for iterative distribution learning) replaces greedy acquisition functions. Yet, there is an opportunity to define new generative workflows that more directly incorporate model-guided optimization methods with consideration of model uncertainty. Batching in MMO can encourage chemical space exploration by rewarding structural diversity, Pareto diversity, or both, but best practices around diversity-aware batching are not well established. Emerging workflows will benefit from the adoption of challenging benchmarks and evaluation metrics that measure the dominated hypervolume or Pareto front density. As newly proposed molecular discovery tools increasingly emphasize multi-objective optimization, emerging methods must address the algorithmic complexities introduced by Pareto optimization.

ABOUT THE AUTHORS

Jenna C. Fromer received her B.S. in Chemical Engineering from Tufts University and is currently a PhD candidate in chemical engineering at MIT. Her research focuses on utilizing multi-objective optimization to accelerate molecular discovery.

Connor W. Coley is the Henri Slezinger (1957) Career Development Assistant Professor of Chemical Engineering and an Assistant Professor of Electrical Engineering and Computer Science at MIT. He received his B.S. from Caltech and PhD from MIT, both in Chemical Engineering. His research group combines expertise in chemical engineering, computer science, and chemistry to improve the utility of computer assistance for chemical discovery.

ACKNOWLEDGMENTS

The authors thank Wenhao Gao, Samuel Goldman, and David Graff for commenting on the manuscript. This work was funded by the DARPA Accelerated Molecular Discovery program under contract HR00111920025.

AUTHOR CONTRIBUTIONS

J.C.F. and C.W.C. jointly prepared the manuscript. J.C.F. prepared all figures.

DECLARATION OF INTERESTS

The authors declare no competing interests.

REFERENCES

- Chong, J.W., Thangalazhy-Gopakumar, S., Muthoosamy, K., and Chemmangattuvalappil, N.G. (2022). Design of bio-oil additives via molecular signature descriptors using a multi-stage computer-aided molecular design framework. *Front. Chem. Sci. Eng.* 16, 168–182.
- Ten, J.Y., Liew, Z.H., Oh, X.Y., Hassim, M.H., and Chemmangattuvalappil, N. (2021). Computer-aided molecular design of optimal sustainable solvent for liquid-liquid extraction. *Process Integr. Optim. Sustain.* 5, 269–284.
- Papadopoulos, A.I., and Linke, P. (2006). Multiobjective molecular design for integrated process-solvent systems synthesis. *AIChE J.* 52, 1057–1070.
- Mah, A.X.Y., Chin, H.H., Neoh, J.Q., Aboagwa, O.A., Thangalazhy-Gopakumar, S., and Chemmangattuvalappil, N.G. (2019). Design of bio-oil additives via computer-aided molecular design tools and phase stability analysis on final blends. *Comput. Chem. Eng.* 123, 257–271.
- Yee, Q.Y., Hassim, M.H., Chemmangattuvalappil, N.G., Ten, J.Y., and Raslan, R. (2022). Optimization of quality, safety and health aspects in personal care product preservative design. *Process Saf. Environ. Prot.* 157, 246–253.
- Ooi, Y.J., Aung, K.N.G., Chong, J.W., Tan, R.R., Aviso, K.B., and Chemmangattuvalappil, N.G. (2022). Design of fragrance molecules using computer-aided molecular design with machine learning. *Comput. Chem. Eng.* 157, 107585.
- Karasuyama, M., Kasugai, H., Tamura, T., and Shitara, K. (2020). Computational design of stable and highly ion-conductive materials using multi-objective bayesian optimization: case studies on diffusion of oxygen and lithium. *Comput. Mater. Sci.* 184, 109927.
- Devereux, L.R., and Cole, J.M. (2021). In *Data Science Applied to Sustainability Analysis*, J. Dunn and P. Balaprakash, eds. (Elsevier), pp. 129–164.
- Hautier, G. (2019). Finding the needle in the haystack: materials discovery and design through computational ab initio high-throughput screening. *Comput. Mater. Sci.* 163, 108–116.
- Hachmann, J., Olivares-Amaya, R., Atahan-Evrenk, S., Amador-Bedolla, C., Sánchez-Carrera, R.S., Gold-Parker, A., Vogt, L., Brockway, A.M., and Aspuru-Guzik, A. (2011). The harvard clean energy project: large-scale computational screening and design of organic photovoltaics on the world community grid. *J. Phys. Chem. Lett.* 2, 2241–2251.
- Ling, J., Hutchinson, M., Antono, E., Paradiso, S., and Meredig, B. (2017). High-dimensional materials and process optimization using data-driven experimental design with well-calibrated uncertainty estimates. *Integr. Mater. Manuf. Innov.* 6, 207–217.
- Jablonka, K.M., Jothiappan, G.M., Wang, S., Smit, B., and Yoo, B. (2021). Bias free multiobjective active learning for materials design and discovery. *Nat. Commun.* 12, 2312.
- Mannodi-Kanakkithodi, A., Pilania, G., Ramprasad, R., Lookman, T., and Gubernatis, J.E. (2016). Multi-objective optimization techniques to design the pareto front of organic dielectric polymers. *Comput. Mater. Sci.* 125, 92–99.
- Hanaoka, K. (2021). Bayesian optimization for goal-oriented multi-objective inverse material design. *iScience* 24, 102781.
- Solomou, A., Zhao, G., Boluki, S., Joy, J.K., Qian, X., Karaman, I., Arróyave, R., and Lagoudas, D.C. (2018). Multi-objective bayesian materials discovery: application on the discovery of precipitation strengthened niti shape memory alloys through micromechanical modeling. *Mater. Des.* 160, 810–827.
- Khatamsaz, D., Vela, B., Singh, P., Johnson, D.D., Allaire, D., and Arróyave, R. (2022). Multi-objective materials bayesian optimization with active learning of design constraints: design of ductile refractory multi-principal-element alloys. *Acta Mater.* 236, 118133.
- Kowalski, J.A., Su, L., Milshtein, J.D., and Brushett, F.R. (2016). Recent advances in molecular engineering of redox active organic molecules for nonaqueous flow batteries. *Curr. Opin. Chem. Eng.* 13, 45–52.
- Winsberg, J., Hagemann, T., Janoschka, T., Hager, M.D., and Schubert, U.S. (2017). Redox-flow batteries: from metals to organic redox-active materials. *Angew. Chem. Int. Ed. Engl.* 56, 686–711.
- Fleitmann, L., Kleinekorte, J., Leonhard, K., and Bardow, A. (2021). Cosmo-Suscampd: sustainable solvents from combining computer-aided molecular and process design with predictive life cycle assessment. *Chem. Eng. Sci.* 245, 116863.
- Wilson, N., and Nimlos, M. (2022). Accelerating materials discovery: artificial intelligence for sustainable, high-performance polymers. <https://www.osti.gov/biblio/1875833>.

21. Melia, H.R., Muckley, E.S., and Saal, J.E. (2021). Materials informatics and sustainability—the case for urgency. *Data-centric Eng.* 2, E19.
22. Raabe, G. (2019). Molecular simulation studies on refrigerants past – present – future. *Fluid Phase Equil.* 485, 190–198.
23. Kazakov, A., McLinden, M.O., and Frenkel, M. (2012). Computational design of new refrigerant fluids based on environmental, safety, and thermodynamic characteristics. *Ind. Eng. Chem. Res.* 51, 537–12548.
24. Sridharan, B., Goel, M., and Priyakumar, U.D. (2022). Modern machine learning for tackling inverse problems in chemistry: molecular design to realization. *Chem. Commun.* 58, 5316–5331.
25. Meyers, J., Fabian, B., and Brown, N. (2021). De novo molecular design and generative models. *Drug Discov. Today* 26, 2707–2715.
26. Segall, M.D. (2012). Multi-parameter optimization: identifying high quality compounds with a balance of properties. *Curr. Pharm. Des.* 18, 1292–1310.
27. Nicolaou, C.A., Brown, N., and Pattichis, C.S. (2007). Molecular optimization using computational multi-objective methods. *Curr. Opin. Drug Discov. Devel.* 10, 316–324.
28. Ekins, S., Honeycutt, J.D., and Metz, J.T. (2010). Evolving molecules using multi-objective optimization: applying to adme/tox. *Drug Discov. Today* 15, 451–460.
29. Nicolaou, C.A., and Brown, N. (2013). Multi-objective optimization methods in drug design. *Drug Discov. Today Technol.* 10, e427–e435.
30. Wei, H., Olarte, M., and Goh, G.B. (2019). Multiple-objective reinforcement learning for inverse design and identification. <http://arxiv.org/abs/1910.03741>.
31. Ramsay, R.R., Popovic-Nikolic, M.R., Nikolic, K., Uliassi, E., and Bolognesi, M.L. (2018). A perspective on multi-target drug discovery and design for complex diseases. *Clin. Transl. Med.* 7, 3.
32. Kim, T., Kim, K., Park, I., Hong, S., and Park, H. (2022). Two-track virtual screening approach to identify the dual inhibitors of wild type and C481S mutant of bruton's tyrosine kinase. *J. Chem. Inf. Model.* 62, 4500–4511.
33. Kuck, D., Singh, N., Lyko, F., and Medina-Franco, J.L. (2010). Novel and selective DNA methyltransferase inhibitors: docking-based virtual screening and experimental evaluation. *Bioorg. Med. Chem.* 18, 822–829.
34. Swaidan, R., Ghanem, B., and Pinnau, I. (2015). Fine-tuned intrinsically ultramicroporous polymers redefine the permeability/selectivity upper bounds of membrane-based air and hydrogen separations. *ACS Macro Lett.* 4, 947–951.
35. Frazier, P.I. (2018). Recent Advances in Optimization and Modeling of Contemporary Problems (INFORMS Tutorials in Operations Research), pp. 255–278.
36. Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., and de Freitas, N. (2016). Taking the human out of the loop: a review of bayesian optimization. *Proc. IEEE* 104, 148–175.
37. Pyzer-Knapp, E.O. (2018). Bayesian optimization for accelerated drug discovery. *IBM J. Res. Dev.* 62, 2:1–2:7.
38. Reker, D., and Schneider, G. (2015). Active-learning strategies in computer-assisted drug discovery. *Drug Discov. Today* 20, 458–465.
39. Ginsbourger, D., Le Riche, R., and Carraro, L. (2010). In *Computational Intelligence in Expensive Optimization Problems*, 2, L.M. Hiot, Y.S. Ong, Y. Tenne, and C.-K. Goh, eds. (Springer Berlin Heidelberg), pp. 131–162.
40. Konakovic Lukovic, M., Tian, Y., and Matusik, W. (2020). Diversity-guided multi-objective bayesian optimization with batch evaluations. *Adv. Neural Inf. Process. Syst.* 17708–17720.
41. Deb, K., and Jain, H. (2014). An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: solving problems with box constraints. *IEEE Trans. Evol. Comput.* 18, 577–601.
42. Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evol. Comput.* 6, 182–197.
43. Verhellen, J. (2022). Graph-based molecular pareto optimisation. *Chem. Sci.* 13, 7526–7535.
44. Agarwal, G., Doan, H.A., Robertson, L.A., Zhang, L., and Assary, R.S. (2021). Discovery of energy storage molecular materials using quantum chemistry-guided multiobjective bayesian optimization. *Chem. Mater.* 33, 8133–8144.
45. Grantham, K., Mukaidaisi, M., Ooi, H.K., Ghaemi, M.S., Tchagang, A., and Li, Y. (2022). Deep evolutionary learning for molecular design. *IEEE Comput. Intell. Mag.* 17, 14–28.
46. González, L.D., and Zavala, V.M. (2022). New paradigms for exploiting parallel experiments in bayesian optimization. Preprint at arXiv. <http://arxiv.org/abs/2210.01071>.
47. Janet, J.P., Ramesh, S., Duan, C., and Kulik, H.J. (2020). Accurate multi-objective design in a space of millions of transition metal complexes with neural-network-driven efficient global optimization. *ACS Cent. Sci.* 6, 513–524.
48. Nicolaou, C.A., Apostolakis, J., and Pattichis, C.S. (2009). De novo drug design using multiobjective evolutionary graphs. *J. Chem. Inf. Model.* 49, 295–307.
49. Zaverkin, V., Holzmüller, D., Steinwart, I., and Kästner, J. (2022). Exploring chemical and conformational spaces by batch mode deep active learning. *Digital Discovery* 1, 605–620.
50. Smith, J.S., Nebgen, B., Lubbers, N., Isayev, O., and Roitberg, A.E. (2018). Less is more: sampling chemical space with active learning. *J. Chem. Phys.* 148, 241733.
51. Kirsch, A., van Amersfoort, J., and Gal, Y. (2019). BatchBALD: efficient and diverse batch acquisition for deep bayesian active learning. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1906.08158>.
52. Li, B., and Rangarajan, S. (2022). A diversity maximizing active learning strategy for graph neural network models of chemical properties. *Mol. Syst. Des. Eng.* 7, 1697–1706.
53. Lipinski, C.A., Lombardo, F., Dominy, B.W., and Feeney, P.J. (2001). Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Adv. Drug Deliv. Rev.* 46, 3–26.
54. Coley, C.W. (2021). Defining and exploring chemical spaces. *Trends Chem.* 3, 133–145.
55. Gómez-Bombarelli, R., Wei, J.N., Duvenaud, D., Hernández-Lobato, J.M., Sánchez-Lengeling, B., Sheberla, D., Aguilera-Iparraguirre, J., Hirzel, T.D., Adams, R.P., and Aspuru-Guzik, A. (2018). Automatic chemical design using a data-driven continuous representation of molecules. *ACS Cent. Sci.* 4, 268–276.
56. Winter, R., Retel, J., Noé, F., Clevert, D.-A., and Steffen, A. (2020). Grüni-fai: interactive multiparameter optimization of molecules in a continuous vector space. *Bioinformatics* 36, 4093–4094.
57. Fu, T., Xiao, C., Li, X., Glass, L.M., and Sun, J. (2021). Mimoso: multi-constraint molecule sampling for molecule optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 125–133.
58. Hartenfeller, M., Proschak, E., Schüller, A., and Schneider, G. (2008). Concept of combinatorial *de novo* design of drug-like molecules by particle swarm optimization. *Chem. Biol. Drug Des.* 72, 16–26.
59. Sowndarya, S.V., Law, J.N., Tripp, C.E., Duplyakin, D., Skordilis, E., Biagioni, D., Paton, R.S., St John, P.C., and St. John, P.C. (2022). Multi-objective goal-directed optimization of *de novo* stable organic radicals for aqueous redox flow batteries. *Nat. Mach. Intell.* 4, 720–730.
60. Ooi, J., Promentilla, M.A.B., Tan, R.R., Ng, D.K., and Chemmangattuvalappil, N.G. (2018). Integration of fuzzy analytic hierarchy process into multi-objective computer aided molecular design. *Comput. Chem. Eng.* 109, 191–202.

61. Liu, X., Chu, J., Zhang, Z., and He, M. (2022). Data-driven multi-objective molecular design of ionic liquid with high generation efficiency on small dataset. *Mater. Des.* 220, 110888.
62. Perron, Q., Mirguet, O., Tajmouati, H., Skiredj, A., Rojas, A., Gohier, A., Ducrot, P., Bourguignon, M.-P., Sansilvestri-Morel, P., Do Huu, N., et al. (2022). Deep generative models for ligand-based de novo design applied to multi-parametric optimization. *J. Comput. Chem.* 43, 692–703.
63. Winter, R., Montanari, F., Steffen, A., Briem, H., Noé, F., and Clevert, D.-A. (2019). Efficient multi-objective molecular optimization in a continuous latent space. *Chem. Sci.* 10, 8016–8024.
64. Urbina, F., Lowden, C.T., Culberson, J.C., and Ekins, S. (2022). MegaSyn: integrating generative molecular design, automated analog designer, and synthetic viability prediction. *ACS Omega* 7, 18699–18713.
65. Firth, N.C., Atrash, B., Brown, N., and Blagg, J. (2015). MOARF, an integrated workflow for multiobjective optimization: implementation, synthesis, and biological evaluation. *J. Chem. Inf. Model.* 55, 1169–1180.
66. Hoffman, S.C., Chenthamarakshan, V., Wadhawan, K., Chen, P.-Y., and Das, P. (2021). Optimizing molecules using efficient queries from property evaluations. *Nat. Mach. Intell.* 4, 21–31.
67. Cardoso Gajo, G., Rodrigues Silva, D., Barigye, S.J., and daCunha, E.F.F. (2018). Multi-objective optimization of benzamide derivatives as rho kinase inhibitors. *Mol. Inform.* 37, 1700080.
68. Chen, Z., Fang, X., Hua, Z., Huang, Y., Wang, F., Wu, H., and Wang, H. (2022). Helix-MO: sample-efficient molecular optimization on scene-sensitive latent space. Preprint at arXiv. <http://arxiv.org/abs/2112.00905>.
69. Jin, W., Barzilay, D., and Jaakkola, T. (2020). Multi-objective molecule generation using interpretable substructures. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 4849–4859.
70. Barshatski, G., Nordon, G., and Radinsky, K. (2021). Multi-property molecular optimization using an integrated poly-cycle architecture. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pp. 3727–3736.
71. Cruz-Monteagudo, M., Borges, F., and Cordeiro, M.N.D.S. (2008). Desirability-based multiobjective optimization for global qsar studies: application to the design of novel nsais with improved analgesic, antiinflammatory, and ulcerogenic profiles. *J. Comput. Chem.* 29, 2445–2459.
72. Cruz-Monteagudo, M., Cordeiro, M.N.D.S., Tejera, E., Dominguez, E.R., and Borges, F. (2012). Desirability-based multi-objective QSAR in drug discovery. *Mini Rev. Med. Chem.* 12, 920–935.
73. Lambrinidis, G., and Tsantili-Kakoulidou, A. (2021). Multi-objective optimization methods in novel drug design. *Expert Opin. Drug Discov.* 16, 647–658.
74. Häse, F., Roch, L.M., and Aspuru-Guzik, A. (2018). Chimera: enabling hierarchy based multi-objective optimization for self-driving laboratories. *Chem. Sci.* 9, 7642–7655.
75. Goel, M., Raghunathan, S., Laghuvarapu, S., and Priyakumar, U.D. (2021). MoleGuLAR: molecule generation using reinforcement learning with alternating rewards. *J. Chem. Inf. Model.* 61, 5815–5826.
76. Pereira, T., Abbasi, M., Oliveira, J.L., Ribeiro, B., and Arrais, J. (2021). Optimizing blood–brain barrier permeation through deep reinforcement learning for de novo drug design. *Bioinformatics* 37, i84–i92.
77. del Rosario, Z., Rupp, M., Kim, Y., Antono, E., and Ling, J. (2020). Assessing the Frontier: active learning, model accuracy, and multi-objective candidate discovery and optimization. *J. Chem. Phys.* 153, 024112.
78. Gopakumar, A.M., Balachandran, P.V., Xue, D., Gubernatis, J.E., and Lookman, T. (2018). Multi-objective optimization for materials discovery via adaptive design. *Sci. Rep.* 8, 3738.
79. Janet, J.P., and Kulik, H.J. (2017). Resolving transition metal chemical space: feature selection for machine learning and structure–property relationships. *J. Phys. Chem. A* 121, 8939–8954.
80. Bilodeau, C., Jin, W., Jaakkola, T., Barzilay, R., and Jensen, K.F. (2022). Generative models for molecular discovery: recent advances and challenges. *WIREs Comput. Mol. Sci.* 12, e1608.
81. Gao, W., and Coley, C.W. (2020). The synthesizability of molecules proposed by generative models. *J. Chem. Inf. Model.* 60, 5714–5723.
82. Sanchez-Lengeling, B., and Aspuru-Guzik, A. (2018). Inverse molecular design using machine learning: generative models for matter engineering. *Science* 361, 360–365.
83. Alshehri, A.S., Gani, R., and You, F. (2020). Deep learning and knowledge-based methods for computer aided molecular design – toward a unified approach: state-of-the-art and future directions. *Comput. Chem. Eng.* 141, 107005.
84. Mouchlis, V.D., Afantitis, A., Serra, A., Fratello, M., Papadiamantis, A.G., Aidinis, V., Lynch, I., Greco, D., and Melagraki, G. (2021). Advances in de novo drug design: from conventional to machine learning methods. *Int. J. Mol. Sci.* 22, 1676.
85. Flam-Shepherd, D., Zhu, K., and Aspuru-Guzik, A. (2022). Language models can learn complex molecular distributions. *Nat. Commun.* 13, 3293.
86. Frey, N.C., Gadepally, V., and Ramsundar, B. (2022). FastFlows: flow-based models for molecular graph generation. Preprint at arXiv. <http://arxiv.org/abs/2201.12419>.
87. Segler, M.H.S., Kogej, T., Tyrchan, C., and Waller, M.P. (2018). Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS Cent. Sci.* 4, 120–131.
88. Boitreau, J., Mallet, V., Oliver, C., and Waldispühl, J. (2020). OptiMol : optimization of binding affinities in chemical space for drug discovery. *J. Chem. Inf. Model.* 60, 5658–5666.
89. Yasonik, J. (2020). Multiobjective de novo drug design with recurrent neural networks and nondominated sorting. *J. Cheminform.* 12, 14.
90. Congreve, M., Carr, R., Murray, C., and Jhoti, H. (2003). A ‘rule of three’ for fragment-based lead discovery? *Drug Discov. Today* 8, 876–877.
91. Iovanac, N.C., MacKnight, R., and Savoie, B.M. (2022). Actively searching: inverse design of novel molecules with simultaneously optimized properties. *J. Phys. Chem. A* 126, 333–340.
92. Abeer, A.N.M.N., Urban, N., Weil, M.R., Alexander, F.J., and Yoon, B.-J. (2022). Multi-objective latent space optimization of generative molecular design models. Preprint at arXiv. <http://arxiv.org/abs/2203.00526>.
93. Venkatasubramanian, V., Chan, K., and Caruthers, J. (1994). Computer-aided molecular design using genetic algorithms. *Comput. Chem. Eng.* 18, 833–844.
94. Sheridan, R.P., and Kearsley, S.K. (1995). Using a genetic algorithm to suggest combinatorial libraries. *J. Chem. Inf. Comput. Sci.* 35, 310–320.
95. Weber, L., Wallbaum, S., Broger, C., and Gubernator, K. (1995). Optimization of the biological activity of combinatorial compound libraries by a genetic algorithm. *Angew. Chem. Int. Ed. Engl.* 34, 2280–2282.
96. Pegg, S.C., Haresco, J.J., and Kuntz, I.D. (2001). A genetic algorithm for structure-based de novo design. *J. Comput. Aided Mol. Des.* 15, 911–933.
97. Brown, N., McKay, B., Gilardoni, F., and Gasteiger, J. (2004). A graph-based genetic algorithm and its application to the multiobjective evolution of median molecules. *J. Chem. Inf. Comput. Sci.* 44, 1079–1087.
98. Jensen, J.H. (2019). A graph-based genetic algorithm and generative model/Monte Carlo tree search for the exploration of chemical space. *Chem. Sci.* 10, 3567–3572.
99. Nigam, A., Pollice, R., and Aspuru-Guzik, A. (2022). Parallel tempered genetic algorithm guided by deep neural networks for inverse molecular design. *Digit. Discov.* 1, 390–404.
100. Durrant, J.D., Lindert, S., and McCammon, J.A. (2013). Autogrow 3.0: an improved algorithm for chemically tractable, semi-automated protein inhibitor design. *J. Mol. Graph. Model.* 44, 104–112.
101. Daeyaert, F., and Deem, M.W. (2017). A pareto algorithm for efficient de novo design of multi-functional molecules. *Mol. Inform.* 36, 1600044.

102. Devi, R.V., Sathya, S.S., and Coumar, M.S. (2021). Multi-objective genetic algorithm for de novo drug design MoGADdrug. *Curr. Comput. Aided Drug Des.* **17**, 445–457.
103. Herring, R.H., and Eden, M.R. (2015). Evolutionary algorithm for de novo molecular design with multi-dimensional constraints. *Comput. Chem. Eng.* **83**, 267–277.
104. Bender, A., and Glen, R.C. (2004). Similarity: a key technique in molecular informatics. *Org. Biomol. Chem.* **2**, 3204–3218.
105. Rosenthal, S. (2020). In *Bioinspired Optimization Methods and Their Applications*, 12438, B. Filipić, E. Minisci, and M. Vasile, eds. (Springer International Publishing), pp. 294–307.
106. Zhou, Z., and Harris, K.D.M. (2008). Counteracting stagnation in genetic algorithm calculations by implementation of a micro genetic algorithm strategy. *Phys. Chem. Chem. Phys.* **10**, 7262–7269.
107. van der Horst, E., Marqués-Gallego, P., Mulder-Krieger, T., van Veldhoven, J., Kruisselbrink, J., Aleman, A., Emmerich, M.T.M., Brussee, J., Bender, A., and IJzerman, A.P. (2012). Multi-objective evolutionary design of adenosine receptor ligands. *J. Chem. Inf. Model.* **52**, 1713–1721.
108. Zhou, Z., Kearnes, S., Li, L., Zare, R.N., and Riley, P. (2019). Optimization of molecules via deep reinforcement learning. *Sci. Rep.* **9**, 10752.
109. Leguy, J., Cauchy, T., Glavatskikh, M., Duval, B., and Da Mota, B. (2020). Evomol: a flexible and interpretable evolutionary algorithm for unbiased de novo molecular generation. *J. Cheminform.* **12**, 55.
110. Khemchandani, Y., O'Hagan, S., Samanta, S., Swainston, N., Roberts, T.J., Bollegala, D., and Kell, D.B. (2020). Deepgraphmolgen, a multi-objective, computational strategy for generating molecules with desirable properties: a graph convolution and reinforcement learning approach. *J. Cheminform.* **12**, 53.
111. Erikawa, D., Yasuo, N., and Sekijima, M. (2021). MERMAID: an open source automated hit-to-lead method based on deep reinforcement learning. *J. Cheminform.* **13**, 94.
112. Horwood, J., and Noutahi, E. (2020). Molecular design in synthetically accessible chemical space via deep reinforcement learning. *ACS Omega* **5**, 32984–32994.
113. Olivecrona, M., Blaschke, T., Engkvist, O., and Chen, H. (2017). Molecular de-novo design through deep reinforcement learning. *J. Cheminform.* **9**, 48.
114. Popova, M., Isayev, O., and Tropsha, A. (2018). Deep reinforcement learning for de novo drug design. *Sci. Adv.* **4**, eaap7885.
115. Pereira, T., Abbasi, M., Ribeiro, B., and Arrais, J.P. (2021). Diversity oriented deep reinforcement learning for targeted molecule generation. *J. Cheminform.* **13**, 21.
116. Neil, D., Segler, M.H., Guasch, L., Ahmed, M., Plumbley, D., Sellwood, M., and Brown, N. (2018). Exploring deep recurrent models with reinforcement learning for molecule design. In 6th International Conference on Learning Representations, Workshop Track Proceedings.
117. Blaschke, T., Arús-Pous, J., Chen, H., Margreitter, C., Tyrchan, C., Engkvist, O., Papadopoulos, K., and Patronov, A. (2020). Reinvent 2.0: an AI tool for de novo drug design. *J. Chem. Inf. Model.* **60**, 5918–5922.
118. Williams, R.J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.* **8**, 229–256.
119. Tan, R.K., Liu, Y., and Xie, L. (2022). Reinforcement learning for systems pharmacology-oriented and personalized drug design. *Expert Opin. Drug Discov.* **17**, 849–863.
120. Bung, N., Krishnan, S.R., and Roy, A. (2022). An in silico explainable multiparameter optimization approach for de novo drug design against proteins from the central nervous system. *J. Chem. Inf. Model.* **62**, 2685–2695.
121. De Cao, N., and Kipf, T.M. (2018). An implicit generative model for small molecular graphs. *icml 2018 workshop on theoretical foundations and applications of deep generative models*. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1805.11973>.
122. Blaschke, T., Olivecrona, M., Engkvist, O., Bajorath, J., and Chen, H. (2018). Application of generative autoencoder in de novo molecular design. *Mol. Inform.* **37**, 1700123.
123. McNaughton, A.D., Bontha, M.S., Knutson, C.R., Pope, J.A., and Kumar, N. (2022). De novo design of protein target specific scaffold-based inhibitors via reinforcement learning. Preprint at arXiv. <http://arxiv.org/abs/2205.10473>.
124. You, J., Liu, B., Ying, R., Pande, V., and Leskovec, J. (2018). Graph convolutional policy network for goal-directed molecular graph generation. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (Red Hook)*, pp. 6412–6422.
125. Ishitani, R., Kataoka, T., and Rikimaru, K. (2022). Molecular design method using a reversible tree representation of chemical compounds and deep reinforcement learning. *J. Chem. Inf. Model.* **62**, 4032–4048.
126. Abbasi, M., Pereira, T., Santos, B.P., Ribeiro, B., and Arrais, J. (2021). Multiobjective reinforcement learning in optimized drug design. 29th european symposium on artificial neural networks. In *Computational Intelligence and Machine Learning*.
127. Liu, X., Ye, K., van Vlijmen, H.W.T., Emmerich, M.T.M., IJzerman, A.P., and van Westen, G.J.P. (2021). Drugex V2: de novo design of drug molecules by pareto-based multi-objective reinforcement learning in polypharmacology. *J. Cheminform.* **13**, 85.
128. Richards, R.J., and Groener, A.M. (2022). Conditional β -VAE for de novo molecular generation. Preprint at arXiv. <http://arxiv.org/abs/2205.01592>.
129. Makhzani, A., Shlens, J., Jaitley, N., Goodfellow, I., and Frey, B. (2016). Adversarial autoencoders. Preprint at arXiv. <http://arxiv.org/abs/1511.05644>.
130. Kang, S., and Cho, K. (2019). Conditional molecular design with deep generative models. *J. Chem. Inf. Model.* **59**, 43–52.
131. Polykovskiy, D., Zhebrak, A., Vetrov, D., Ivanenkov, Y., Aladinskiy, V., Mamoshina, P., Bozdaganyan, M., Aliper, A., Zhavoronkov, A., and Kadurin, A. (2018). Entangled conditional adversarial autoencoder for de novo drug discovery. *Mol. Pharm.* **15**, 4398–4405.
132. Simonovsky, M., and Komodakis, N. (2018). GraphVAE: towards generation of small graphs using variational autoencoders. In *Artificial Neural Networks and Machine Learning – ICANN 2018*, pp. 412–422.
133. Lim, J., Ryu, S., Kim, J.W., and Kim, W.Y. (2018). Molecular generative model based on conditional variational autoencoder for de novo molecular design. *J. Cheminform.* **10**, 31.
134. Lee, M., and Min, K. (2022). Multi-objective inverse design via molecular graph conditional variational autoencoder. *J. Chem. Inf. Model.* **62**, 2943–2950.
135. Leach, A.G., Jones, H.D., Cosgrove, D.A., Kenny, P.W., Ruston, L., MacFaul, P., Wood, J.M., Colclough, N., and Law, B. (2006). Matched molecular pairs as a guide in the optimization of pharmaceutical properties; a study of aqueous solubility, plasma protein binding and oral exposure. *J. Med. Chem.* **49**, 6672–6682.
136. Tyrchan, C., and Evertsson, E. (2017). Matched molecular pair analysis in short: algorithms, applications and limitations. *Comput. Struct. Biotechnol. J.* **15**, 86–90.
137. Jin, W., Yang, K., Barzilay, R., and Jaakkola, T. (2019). Learning multi-modal graph-to-graph translation for molecular optimization. In *International Conference on Learning Representations*.
138. Wang, Z., Nie, W., Qiao, Z., Xiao, C., Baraniuk, R., and Anandkumar, A. (2022). Retrieval-based controllable molecule generation. Preprint at arXiv. <http://arxiv.org/abs/2208.11126>.
139. Irwin, R., Dimitriadis, S., He, J., and Bjerrum, E.J. (2022). Chemformer: a pre-trained transformer for computational chemistry. *Mach. Learn. Sci. Technol.* **3**, 015022.
140. He, J., Nittinger, E., Tyrchan, C., Czechtizky, W., Patronov, A., Bjerrum, E.J., and Engkvist, O. (2022). Transformer-based molecular optimization beyond matched molecular pairs. *J. Cheminform.* **14**, 18.

141. Bagal, V., Aggarwal, R., Vinod, P.K., and Priyakumar, U.D. (2022). MolGPT: molecular generation using a transformer-decoder model. *J. Chem. Inf. Model.* 62, 2064–2076.
142. Kotsias, P.-C., Arús-Pous, J., Chen, H., Engkvist, O., Tyrchan, C., and Bjerrum, E.J. (2020). Direct steering of de novo molecular generation with descriptor conditional recurrent neural networks. *Nat. Mach. Intell.* 2, 254–265.
143. He, J., You, H., Sandström, E., Nittinger, E., Bjerrum, E.J., Tyrchan, C., Czechtizky, W., and Engkvist, O. (2021). Molecular optimization by capturing chemist's intuition using deep neural networks. *J. Cheminform.* 13, 26.
144. Mukaidaisi, M., Vu, A., Grantham, K., Tchagang, A., and Li, Y. (2022). Multi-objective drug design based on graph-fragment molecular representation and deep evolutionary learning. *Front. Pharmacol.* 13, 920747.
145. Li, Y., Zhang, L., and Liu, Z. (2018). Multi-objective de novo drug design with conditional graph generative model. *J. Cheminform.* 10, 33.
146. Domenico, A., Nicola, G., Daniela, T., Fulvio, C., Nicola, A., and Orazio, N. (2020). *De novo* drug design of targeted chemical libraries based on artificial intelligence and pair-based multiobjective optimization. *J. Chem. Inf. Model.* 60, 4582–4593.
147. Wang, J., Hsieh, C.-Y., Wang, M., Wang, X., Wu, Z., Jiang, D., Liao, B., Zhang, X., Yang, B., He, Q., et al. (2021). Multi-constraint molecular generation based on conditional transformer, knowledge distillation and reinforcement learning. *Nat. Mach. Intell.* 3, 914–922.
148. Ståhl, N., Falkman, G., Karlsson, A., Mathiason, G., and Boström, J. (2019). Deep reinforcement learning for multiparameter optimization in *de novo* drug design. *J. Chem. Inf. Model.* 59, 3166–3176.
149. Chen, B., Fu, X., Barzilay, R., and Jaakkola, T. (2021). Fragment-based sequential translation for molecular optimization. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2111.01009>.
150. Collette, Y., and Siarry, P. (2005). Three new metrics to measure the convergence of metaheuristics towards the pareto Frontier and the aesthetic of a set of solutions in biobjective optimization. *Comput. Oper. Res.* 32, 773–792.
151. Li, M., and Yao, X. (2020). Quality evaluation of solution sets in multiobjective optimisation: a survey. *ACM Comput. Surv.* 52, 1–38.
152. Maltese, J., Ombuki-Berman, B.M., and Engelbrecht, A.P. (2018). A scalability study of many-objective optimization algorithms. *IEEE Trans. Evol. Comput.* 22, 79–96.
153. Renz, P., Van Rompaey, D., Wegner, J.K., Hochreiter, S., and Klambauer, G. (2019). On failure modes in molecule generation and optimization. *Drug Discov. Today Technol.* 32–33, 55–63.
154. Gao, W., Fu, T., Sun, J., and Coley, C.W. (2022). Sample efficiency matters: a benchmark for practical molecular optimization. Preprint at arXiv. <http://arxiv.org/abs/2206.12411>.
155. Ryan, E.G., Drovandi, C.C., McGree, J.M., and Pettitt, A.N. (2016). A review of modern computational algorithms for bayesian optimal design. *Int. Stat. Rev.* 84, 128–154.