



Universiteit
Leiden
The Netherlands

Artificial intelligence in multi-objective drug design

Luukkonen, S.I.M.; Maagdenberg, H.W. van den; Emmerich, M.T.M.; Westen, G.J.P. van

Citation

Luukkonen, S. I. M., Maagdenberg, H. W. van den, Emmerich, M. T. M., & Westen, G. J. P. van. (2023). Artificial intelligence in multi-objective drug design. *Current Opinion In Structural Biology*, 79. doi:10.1016/j.sbi.2023.102537

Version: Publisher's Version

License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)

Downloaded from: <https://hdl.handle.net/1887/3590043>

Note: To cite this publication please use the final published version (if applicable).



Artificial intelligence in multi-objective drug design

Sohvi Luukkonen^{1,a}, Helle W. van den Maagdenberg^{1,a},
 Michael T. M. Emmerich² and Gerard J. P. van Westen¹

Abstract

The factors determining a drug's success are manifold, making *de novo* drug design an inherently multi-objective optimisation (MOO) problem. With the advent of machine learning and optimisation methods, the field of multi-objective compound design has seen a rapid increase in developments and applications. Population-based metaheuristics and deep reinforcement learning are the most commonly used artificial intelligence methods in the field, but recently conditional learning methods are gaining popularity. The former approaches are coupled with a MOO strategy which is most commonly an aggregation function, but Pareto-based strategies are widespread too. Besides these and conditional learning, various innovative approaches to tackle MOO in drug design have been proposed. Here we provide a brief overview of the field and the latest innovations.

Addresses

¹ Leiden Academic Centre for Drug Research, Leiden University, Einsteinweg 55, Leiden, 2333 CC, the Netherlands

² Leiden Institute of Advanced Computer Science, Leiden University, Niels Bohrweg 1, Leiden, 2333 CC, the Netherlands

Corresponding author: van Westen, Gerard J.P. (gerard@iacdr.leidenuniv.nl)

 (Luukkonen S.)

 (van Westen G.J.P.)

^a Authors have contributed equally to this paper.

Current Opinion in Structural Biology 2023, 79:102537

This review comes from a themed issue on **Artificial Intelligence (AI) Methodology in Structural Biology (2023)**

Edited by **Andreas Bender, Chris de Graaf and Noel O'Boyle**

For complete overview of the section, please refer the article collection - [Artificial Intelligence \(AI\) Methodology in Structural Biology \(2023\)](#)

Available online 10 February 2023

<https://doi.org/10.1016/j.sbi.2023.102537>

0959-440X/© 2023 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Keywords

Multi-objective optimisation, Pareto dominance, *de novo* drug design, Compound optimisation.

Introduction

Drug discovery is a challenging and expensive process. On average, taking more than ten years and costing more

than two billion dollars [1]. Computer-aided drug design can reduce these costs by early discontinuation of failed compounds and reducing the number of experiments needed. In *de novo* drug design (DNDD), the vast space of drug-like molecules ($\sim 10^{63}$) [2] is explored in the search for attractive novel drug candidates. Typically the predicted efficacy, synthesizability, and drug-likeness of the compounds need to be maximised while off-target effects and toxicity minimised. In addition, favourable physicochemical and pharmacokinetic properties need to be obtained. This makes DNDD inherently a multi-objective optimisation (MOO) problem.

The field of MOO seeks to develop algorithmic methods for simultaneously minimising (or maximising) multiple objectives. Due to conflicting objective functions, in MOO, there typically exists no single best solution, but one needs to choose out of the following three approaches [3].

- *A priori*: Aggregate the objective functions into a single objective function employing a predefined scalarisation function, which is then optimised.
- *A posteriori*: Determine or approximate the set of non-dominated solutions (Pareto optimal set). After non-dominated ranking, the trade-off between solutions is assessed.
- *Progressive or interactive*: Alternate between automated search and preference elicitation phases, in which the decision maker refines preferences (e.g., weights or regions of interest) based on inspection of intermediate results.

The advent of machine learning-based prediction tools and optimisation algorithms that could handle complex, non-numerical solution representations (e.g., chemical graphs) enabled the application of MOO in molecular design. Such algorithms typically belong to the class of population-based metaheuristics, among which multi-objective evolutionary algorithms and (particle) swarm optimisation algorithms are major subclasses [4,5]. Starting with an initial set of compounds, these methods create new molecules by structural modifications at each iteration and apply selection operators that favour Pareto optimal solutions and provide diversity [6,7].

In the last decade, these methods have been accompanied by generative methods that sample candidate molecules in promising parts of the chemical space and construction heuristics that, to generate candidate solutions, start with an empty compound description (but not necessarily without prior knowledge), and extend it step-by-step (*e.g.* by adding groups or atoms, deciding based on learned probabilities for good ‘moves’). Typically these generative methods are optimised with transfer (TL) and reinforcement learning (RL) [8]. Recently multiple models with conditional learning (CL), where the desired ranges of properties are defined *a priori* and passed as additional input to the generator model, have emerged.

A plethora of published molecular generators with MOO is available; examples here include recurrent neural networks (RNNs), generative adversarial networks (GANs), graph convolutional policy networks (GCPN) with MOO in a reinforcement learning framework, but also conditional variational autoencoders (cVAEs), RNNs (cRNNs), transformers (cTrans) and genetic algorithms (GAs) among others. For a recent review, see Liu et al. [9].

These generators are coupled with different formulations of the objective functions. The variety of molecule scoring approaches is illustrated in Figure 1. Early applications of MOO often used physicochemical properties to generate drug-like molecules. Due to the enormous progress in modern AI methods, there has been a major increase in the use of ML-based quantitative structure–activity relationship (QSAR) models for generating target-specific compounds. Evolutionary and swarm optimisation algorithms use the objective function(s) as selection operator(s), while in reinforcement and conditional learning they are part of the reinforcement and conditional learning optimisation scheme.


In section 2, we briefly introduce classic MOO approaches, followed by recent applications and developments in DNDD in section 3, before concluding with final remarks on the perspectives in the field.

Methods of MOO

A priori approach - aggregation methods

A straightforward and most commonly used method to deal with the complexities of MOO is to convert the MOO to single-objective optimisation. A scalarisation function is employed to aggregate all the objectives into a single scalar objective function called the multi-objective desirability function. Frequently used scalarisation functions include the weighted arithmetic and geometric means or the Chebychev scalarisation. If the sum of the weight equals 1, weighted arithmetic and geometric means are reduced to the weighted sum (WS)

Figure 1



Molecule scoring approaches used in objective functions: A human expert performing the selection [10] (blue), quick to calculate physico-chemical properties of drug-like molecules (green), structure-based molecule scoring (yellow), machine learning-based QSAR predictions (orange), or so-called closed-loop drug discovery where iteratively molecules are generated, synthesised, tested and experimental information is fed back to the model (purple) [11]. The thin arrows indicate the flow of information between the different scoring methods.

and product (WP), respectively. In Figure 2a-c., we give their formulas and illustrate decomposition ranking on a two-objective example. A common optimisation task is maximising the similarity to an existing compound and the QED. Here we illustrate different rankings by maximising similarity to methotrexate, a chemotherapy agent and immune-system suppressant.

As illustrated by comparing rankings in Figure 2a-c To Figure 2d., a multi-objective desirability function is not guaranteed to find the full optimal Pareto front, since the form of the function can strongly affect the ranking of the molecules. Typically, scalarization methods only result in a single best solution. In the case of objective maximisation, the weighted sum (Figure 2b.) will reward compounds that perform well on average, whereas the Chebychev scalarisation, with the origin as a reference point (Figure 2c.), will reward compounds that perform well on the best performing objective. The weighted product (Figure 2a.) performs well when the least performing objective performs well. All individual objectives need to be normalised to the same range. Modifier functions are often used to clip all raw single-objective values between 0 (undesirable) and 1 (desirable) before combining.

Figure 2




Illustration of two-objective ranking of 1000 example molecules based on maximisation of the QED and Tanimoto similarity to methotrexate. Top: Ranking with aggregation functions: A. weighted geometric mean (WGM), B. weighted arithmetic mean (WAM), and C. weighted Chebychev scalarisation (WCS), and below their formulas where w_i and x_i are the weight and value of the i th objective and f a modifier function. In these examples $w_{QED} = w_{Tsim} = 0.5$ and $f(QED) = QED$ and $f(T_{sim}) = \text{MinMax}(T_{sim})$. Bottom: D. Pareto ranking with Pareto front marked with dark blue hexagons, and example molecules from the Pareto front.

A posteriori approach - Pareto-based methods

In contrast to aggregation, Pareto-based optimisation methods do not combine multiple objectives into one but rather search for the best trade-off between them. In Pareto ranking (PR), one solution is only considered better than another when it is better or equal in every objective and better in at least one objective. This means there are several equivalent best or so-called 'non-dominated' solutions. The set of optimal solutions in the objective space forms the Pareto front (Figure 2d).

One of the most widely used Pareto ranking (PR) methods is the Non-dominated Sorting Genetic Algorithm (NSGA-II) [12]. It is a genetic algorithm where the new generation is selected from the current population by ranking the parent and new generation of solutions through applying non-dominated sorting and subsequently sorting by crowding distance to increase the diversity of solutions.

In general, the convergence time of Pareto-frontier computation is higher than those of scalarization

methods, because in the first case the computational resources are distributed over the points (compounds) of an approximation set, whereas in the latter case only a single point (compound) has to be improved. In case of many objectives, also the ranking by an aggregation method is faster and scales slightly better than with a method using Pareto frontiers, but often in DNDD the difference in computation time is small compared to the other steps: the molecules generation and scoring them on different objectives.¹ The scalarisation functions are more stable when moving to a higher number of objectives, as the number of equivalent Pareto solutions increases with the number of dimensions. On the other hand, the Pareto-based schemes enforce diversity in the ranking which is not the case for the multi-objective desirability functions. To circumvent the lack of diversity and the risk of being stuck in local optima

¹ The weighted sum - and the NSGA-II ranking procedure have a worst-case complexity of $O(MN \log N)$ and $O(\min(MN^2, N(\log N)^{(M-1)}))$ (Jensen-Fortin-Buzdalov algorithm), respectively, with M the number of objectives and N the number of compounds.

aggregation methods can be coupled with the diversity filters [13].

Recent developments and applications in DNDD

In Table 1, we summarise the recent developments and application of MOO in DNDD with and without RL, and in the following sub-sections, we present recent developments and applications of MOO in DNDD by optimisation strategy.

Population-based metaheuristics

The versatility of nature-inspired algorithms makes them one of the most widely used methods for molecular generation. With LigBuilder v3, Yuan et al. [14] present a structure-based DNDD method to design ligands to target multiple receptors, binding sites, or various conformers. They propose GA-based strategies to build molecules into the multiple binding site structures and combine the binding free energies in their desirability function. Other recently used nature-based MOO algorithms with an aggregation function include works by Bilsland et al. [15] with swarm optimisation algorithms applied using a ‘dual’ autoencoder, encoding and decoding SMILES and molecular fingerprints simultaneously, and SAGS [16] which uses simulated annealing with a graph neural network to optimise molecular graphs.

The use of NSGA-II for DNDD is exemplified in Ref. [17] with the SELFIES molecular representation. SELFIES, unlike SMILES, are always syntactically correct; therefore, mutation operators can be used without creating non-valid molecules. Verhellen [18] combined a graph-based representation with NSGA-II and NSGA-III, which uses so-called reference directions to select diverse compounds instead of the crowding distance. In Ref. [19], the authors propose a variation of LigBuilder where they replace the GA-strategy with their own multi-objective bio-film optimisation algorithm, MoBifi; a MOO algorithm that mimics the behaviour of bacteria in bio-film using Pareto dominance.

Reinforcement learning

Due to the prominence of RNN-based natural language processing models in the last decade, a plethora of SMILES-based RNN models with RL framework are still being developed and applied [20–29]. The RL framework has also been used with other generator architectures such as GCPNs [30,31], GANs [32,33], and cTrans [34,35].

A large majority of these methods employ a WS or WP scalarisation function with weights that are fixed hyperparameters [20–24,30,34–36]; the most prominent being REINVENT 2.0 [20], a production-ready DNDD

package that has become a reference tool in the field and the base algorithm for subsequent developments and extensions [21,37,25,38]. In their work, Perron et al. [24] use a WS composed of 13 objectives, affinities, and ADME properties in a practical drug discovery project, claiming added value from a DNDD with a MOO approach over traditional medicinal chemistry approaches. Out of the 11 top predicted compounds they synthesised and tested, 3 fulfilled all 13 objectives. In a different approach, DrugEx v2 [27] and DeepFMPO v3D [28] use a parametric-WS. It uses dynamic weights to reward compounds that score well in objectives where most molecules are performing badly. At each iteration these weights are updated.

As for the use of PR methods, DrugEx v2 [27] proposes a novel variation of NSGA-II to rank the molecules in the same Pareto front using the Tanimoto distance between molecules instead of the crowding distance. Abbasi et al. [33] use Fonseca and Fleming’s Due to the prominence sorting algorithm, where dominated molecules are scored based on how many predicted molecules dominate them. Alberga et al. present pair-wise Pareto optimisation in Ref. [38], a novel PR approach. Instead of applying PR on all objectives simultaneously, the solutions are ranked using each combination of two objectives. The authors claim that this method provides an advantage over traditional PR with an increasing number of objectives as the number of equivalent solutions will also increase in this scenario.

Both Goel et al. [26] and Guo et al. [25] propose using a sequence of single objectives to tackle the MOO problem. In Guo et al. [25], the objectives are arranged as a sequence of increasing complexity and the agent is trained on a single objective until convergence before moving to the next objective. This is shown to accelerate agent convergence on complex MOO when compared to the baseline RL. Whereas in MoCuLAR [26], the agent is optimised with alternating rewards where the objective is changed every n iteration.

RationaleRL [36] is a novel approach using ‘rationales’ which are molecular substructures that are linked to the desired property in a compound. First, rationales per objective are extracted from a set of molecules with the desired property by pruning molecular graphs while maintaining the desired property by a Monte Carlo Tree Search algorithm. Then rationals for multiple objectives are combined and completed using a VAE in a RL framework.

Conditional learning

In recent years, there has been a strong interest in conditional learning for molecular generation. These conditional generators show promising results for MOO of compounds and are a serious alternative to the

Table 1

Recent *de novo* drug design methods with multi-objective optimization.

Method	Ref.	Architecture ^a	Representation	MOO Strategy ^b	Objectives ^c	Open Source
GNC	[45]	AE + DNN	continuous	WS	[3] Affinities (BACE1, ALK or CDK4 and CDK6) and TS	No
<i>Bilsland</i>	[15]	AE + PSO	SMILES + fingerprints	WS	[6] fragment score, SA, heavy atom count, FSP3, and undesirable functionalities	Yes
<i>Kotsias</i>	[39]	cRNN	SMILES	CL	[7] Affinity (DRD2), logP, TPSA, MW, QED, HBA and HBD	Yes
CMG	[42]	cTrans	SMILES	CL	[4] Affinity (DRD2), logP, QED and TS	Yes
MolGPT	[43]	cTrans	SMILES	CL	[4] logP, SA, TPSA and QED	Yes
<i>He</i>	[49]	cTrans	SMILES	CL	[3] logP, solubility and clearance	Yes
MCMG	[35]	cTrans + RL	SMILES	CL + WS	[4] Affinities (DRD2 or JNK3 and GSJ3 β), QED and SA	Yes
GCT	[41]	cTrans + VAE	SMILES	CL	[3] logP, TPSA and QED	Yes
MGCVAE	[40]	cVAE	graph	CL	[2] logP and molar refractivity	Yes
PaccMann ^{RL}	[34]	cVAE + RL	SELFIES	CL + WS	[2] Affinity (41 SARS-CoV-2 targets) and toxicity	Yes
LigBuilder v3	[14]	GA	3D	WS	Ligand affinity and efficiency (HIV protease and HIV), drug-likeness and MCF	Yes
<i>Elend</i>	[50]	GA	SMILES	WS	[5] DS (SARS-CoV-2 main protease), SA, QED, natural product-likeness and toxicity	No
<i>Cofala</i>	[17]	GA	SELFIES	WS/PR	[5] DS (SARS-CoV-2 main protease), SA, QED, natural product-likeness and toxicity	No
<i>Verhellen</i>	[18]	GA	graph	PR	[5] Affinities (hERG, SCN2A, DAPk1, DRP1, ZIPK, 5-HT2A, 5-HT2B and DR2D) or sets of 5 GuacaMol tasks	Yes
MoBifi	[19]	GA	SELFIES	PR	[3] TS, oral bioavailability, Veber score	No
DLGN	[32]	GAN + RL	SMILES	WS	[2] Affinities (DRD2 and HTR1A)	Yes
<i>Abbasi</i>	[33]	GAN + RL	continuous	GAN + PR	[2] Affinity (ADORA2A or KOR) and logP, SA or TPSA	Yes
DeepGraphMolGen	[30]	GCPN + RL	graph	WS	[2] Affinities (Dopamine and norepinephrine transporters)	Yes
MNCE-RL	[31]	GCPN + RL	graph	WS/WP	[2] TS, logP or QED	Yes
SAGS	[16]	GNN + SA	graph/SMILES	WGM	[2] logP and QED	Yes
<i>Iovanac</i>	[48]	gVAE + TL	Grammar parse trees based on SMILES	CL + Active learning	[3] vertical ionization potential, electron affinity, dipole moment	Yes
RationalIRL	[36]	MCTS + VAE + RL	graph	WS	[4] Affinities (JNK3 and GSJ3 β), QED and SA	Yes
STONED	[46]	mutations	SELFIES	median molecules	[3] LUMO, dipole moment, HOMO–LUMO gap	Yes
REACTOR	[51]	RL	SMARTS	WCS	[4] Affinities (DRD1, DRD2, DRD3), logP, MW, and absorption	no
Megasyn	[22]	RNN + RL	SMILES	WS	[6] Affinities (HER1, HER2 and HERG), QED, TS and BBB	On request
<i>Bung</i>	[23]	RNN + RL	SMILES	WS	[4] DS (5-HT1B), MW, BBB and logP	On request
<i>Perron</i>	[24]	RNN + RL	SMILES	WS	[13] 7 Affinities (5-HT2A, 5-HT2B, alpha1, D1, Nav1.2, hERG, 1 undisclosed), 4 ADME	No

(continued on next page)

Table 1. (continued)

Method	Ref.	Architecture ^a	Representation	MOO Strategy ^b	Objectives ^c	Open Source
REINVENT 2.0	[20]	RNN + RL	SMILES	WS/WP	assays (microsomal stability on human (HLM) and rat (RLM) and permeability and efflux Caco2 assays) and TS and QED	Yes
LibINVENT	[21]	RNN + RL	SMILES	WS/WP	–	Yes
<i>Pereira</i>	[29]	RNN + RL	SMILES	WS/WCS	[2] Affinity (DRD2) or ROCS and reaction filters	Yes
DeepFMPO v3D	[28]	RNN + RL	SMILES + 3D	pWS	[2] Affinity (ADORA2A) and BBB	Yes
DrugEx v2	[27]	RNN + RL	SMILES	pWS/PR	[3] MW, logP and TPSA	Yes
<i>Guo</i>	[25]	RNN + RL	SMILES	Curriculum learning	[3] Affinities (ADORA2A, ADORA2B, hERG)	Yes
<i>Alberga</i>	[38]	RNN + RL	SMILES	WS + pair-wise PR	[3] DS (PDK1), TS, ROCS and QED	Yes
MoleGuLAR	[26]	RNN + RL	SMILES	Alternative rewards	[5] Affinities (NA, AChe or SARS-CoV-2 M _{pro}) and subsets of MW, logP, HBD, HBA, aliphatic rings, TS	Yes
<i>Yasonik</i>	[47]	RNN + TL	SMILES	PR	[5] DS (TTBK1 or SARS-CoV-2 M _{pro}), dG _{hyd} , QED, logP and TPSA	Yes
					[5] logP, MW, HBA, HBD, rotatable bonds	Yes

^a AE - autoencoder, DNN - deep neural network, GA - genetic algorithm, GAN - generative adversarial network, GNN - graph neural network, GCPN - graph convolutional policy network, PSO - particle swarm optimization, RL - reinforcement learning, TL - transfer learning, (c)RNN - (conditional) recurrent neural network, SA - simulated annealing, (c)Trans - (conditional) transformer, (c)VAE - (conditional) variational autoencoder.

^b CL - conditional learning, GAN - generative adversarial network, PR - Pareto ranking, WCS - weighted Chebychev scalarization, WGM - weighted geometric mean, WP - weighted product, (p)WS - (parametric) weighted sum.

^c BBB - blood–brain barrier, DS - docking score, CNS - central nervous system desirability score, FSP3 - fraction sp³-hybridized carbons, HBA/B–hydrogen bond acceptors/donors, (p)logP - (penalized) partition coefficient, MCF - medical chemical filter of toxic fragments, MW - molecular weight, QED - quantitative estimation of drug-likeness, SA - synthetic accessibility, TPSA - topological polar surface area, TS - Tanimoto similarity.

computationally expensive optimisation loops of RL with MOO. Kotsias *et al.* [39] proposed a SMILES-based cRNN, Lee *et al.* [40] a molecular graph cVAE, and a variety of conditional transformer models for DNDD have been proposed [34,35,41–44]. The cRNN differs from the other models as the generation is conditioned with molecular properties only without molecular encoding during training or generation.

In 2021, two methods combining conditional generators with the RL framework were proposed. PacMann^{RL} [34] uses a cTrans combined with a distillation model to create the prior for the RL instead of using a transfer learning step. In MCMG [35], the RL agent is obtained by first training two VAEs, one for SELFIES and another one for proteins, and then combining the pretrained protein-encoder with the pretrained molecule-decoder.

Other methods

Gao *et al.* [45] propose an original approach, a Generative Network Complex, where an input molecule is encoded into latent space and then optimised with a DNN by minimising the multi-objective loss with gradient descent. Another quite distinct method is proposed by Nigam *et al.* [46]. To find molecules optimised for multiple properties, a molecule scoring high per objective is selected and then median molecules that are chemically close to the target molecules are sought. To find these median molecules, they generate the local chemical subspace around molecules using mutations of multiple SELFIES representing the target molecule.

Yasonik *et al.* [47] use the evolutionary approach to optimise a RNN. Generated compounds are ranked with Fonseca and Fleming's non-dominated sorting algorithm, and the top compounds are selected and used to fine-tune the agent through TL. In their work, Iovanac *et al.* [48] use an active learning process to improve their grammar VAE coupled with linear property predictors using the later features. The latent features are conditionally sampled during generation to fulfill given property ranges and decoded to SMILES. The generated molecules are then inspected with a quantum calculation and the model is retrained with promising compounds.

The main advantage of population-based metaheuristics is that they do not require a huge initial dataset to create drug-like molecules in contrast to pre-training generative models. They are also faster to train than generative models with reinforcement or transfer learning, but their efficiency depends strongly on the initial population. Furthermore, population-based models only learn optimal features and information about unwanted features is lost at each iteration. The generative models are able to learn, in addition to the optimal features, unwanted ones and avoid creating them. Population-based

metaheuristics and reinforcement learning require looping over computationally expensive scoring, which is not the case for conditional learning.

Concluding remarks

In the past few years, a tremendous number of studies have been published on the topic of multi-objective compound optimisation; among these were novel methods as well as techniques adapted from other disciplines. The vast majority of new *de novo* drug design applications incorporate MOO, and a gradual shift towards an increasing number of objectives is occurring with the development of new methods. This progression towards many-objective optimisation (MaOO, more than three objectives) is very relevant for drug discovery as not few but many factors determine the eventual success of a drug. Although currently, MaOO is mostly applied with weighted sums, an example of an explicit many-objective algorithm, NSGA-III [18], has recently been applied to compound optimisation. There is a multitude of many-objective algorithms applied in other domains that could be adopted for drug discovery [52].

Even though reinforcement learning dominates the field of DNDD with MOO, mainly coupled with SMILES-based RNNs, some older approaches persist (*e.g.*, GAs, PSO, and SA), and new designs emanate in model architecture as well as in molecular representation. Especially conditional models have emerged as a powerful tool for multi-objective compound design and their potential for MaOO should be further investigated.

The consolidation of the multi-objective compound optimisation field leads to the increased need for benchmarking. Many recent papers use Guacamol benchmarks [53], MOSES [54] or compare their models to state-of-the-art methods. Reeves *et al.* [55] also provide a framework to structurally compare the different elements of *de novo* drug design algorithms. However, there seems to be no real consensus on best practices. Moreover, we noted that the code and data are often available, but very few models have been experimentally validated so far. Of the recent applications discussed in this review, only one included experimental validation [24] and two computational validation with molecular dynamics simulations [14,50]. From our perspective, the field would benefit strongly from standardisation of method benchmarking and, perhaps most important, although difficult and time-consuming, experimental validation.

The importance of synthesizability and retrosynthetic planning in DNDD is often noted and sometimes done in a post-processing step, *e.g.*, in Refs. [14,34,50,15,16,22], or rarely included as part of the MOO during training and generation beyond the very approximate synthetic accessibility score [51]. We expect recent advances in

ML-accelerated retrosynthetic design [56] to enable its use in the MOO process in the future.

Other important developments in the broader field of AI research are explainable AI [57] and the uncertainty quantification of predictions which could be used in bayesian multi-objective optimisation [58,59]. With the increasing complexity and adoption of ML, there is a growing need for more interpretable and transparent models. In the context of chemoinformatics and DNDD, it is essential that medicinal chemists are able to trust and understand the model predictions so they will consider the results in their decision-making process. Only a few of the articles discussed here paid attention to the explainability of their models. In Ref. [36], this issue is addressed by initially identifying sub-structures linked to desired properties before using reinforcement learning to combine these into novel compounds. The importance of uncertainty quantification in decision-making and predictor quality in the DNDD process has been noted, but is yet to be addressed by the community [60]. Especially when exploring new chemical space, there is a risk of falling outside the domain of applicability of the predictors. Further research in QSAR modelling with uncertainty estimation and how to incorporate them in the MOO should be done to improve multi-objective compound design.

Declaration of competing interest

Nothing declared.

Data availability

No data was used for the research described in the article.

Acknowledgments

SL acknowledges the Dutch Research Council (NWO ENPPS. LIFT.019.010) for funding.

References

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

* of special interest

- DiMasi JA, Grabowski HG, Hansen RW: **Innovation in the pharmaceutical industry: new estimates of R&D costs.** *J Health Econ* 2016, **47**:20–33, <https://doi.org/10.1016/j.jhealeco.2016.01.012>.
- Kirkpatrick P, Ellis C: **Chemical space.** *Nature* 2004, **432**: 823–824, <https://doi.org/10.1038/432823a>.
- Miettinen K: *Nonlinear Multiobjective optimization, Vol. 12 of International Series in Operations research Management Science.* Springer Science & Business Media; 1999, <https://doi.org/10.1007/978-1-4615-5563-6>.
- Deb K: *Multi-objective optimisation using evolutionary algorithms.* Wiley Interscience Series in Systems and Optimization, John Wiley & Sons; 2001. Third print 2004).
- Emmerich M, Deutz AH: **A tutorial on multiobjective optimization: fundamentals and evolutionary methods.** *Nat Comput* 2018, **17**:585–609, <https://doi.org/10.1007/s11047-018-9685-y>.
- Nicolaou CA, Brown N: **Multi-objective optimization methods in drug design.** *Drug Discov Today Technol* 2013, **10**:e427–e435, <https://doi.org/10.1016/j.ddtec.2013.02.001>.
- Kruisselbrink JW, Aleman A, Emmerich MT, IJzerman AP, Bender A, Bäck T, van der Horst E: **Enhancing search space diversity in multi-objective evolutionary drug molecule design using niching.** In *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*; 2009:217–224, <https://doi.org/10.1145/1569901.1569932>.
- Elbadawi M, Gaisford S, Basit AW: **Advanced machine-learning techniques in drug discovery.** *Drug Discov Today* 2021, **26**: 769–777, <https://doi.org/10.1016/j.drudis.2020.12.003>.
- Liu X, IJzerman AP, van Westen GJP: **Computational approaches for de novo drug design: past, present, and future.** In *H. Cartwright (Ed.), Artificial Neural Networks, Methods in Molecular Biology.* New York, NY: Springer US; 2021:139–165, https://doi.org/10.1007/978-1-0716-0826-5_6.
- Lameijer E-W, Kok JN, Bäck T, IJzerman AP: **The molecule evaluator. an interactive evolutionary algorithm for the design of drug-like molecules.** *J Chem Inf Model* 2006, **46**: 545–552, <https://doi.org/10.1021/ci050369d>.
- Saikin SK, Kreisbeck C, Sheberla D, Becker JS, Aspuru-Guzik A: **Closed-loop discovery platform integration is needed for artificial intelligence to make an impact in drug discovery.** *Expert Opin Drug Discov* 2019, **14**:1–4, <https://doi.org/10.1080/17460441.2019.1546690>.
- Deb K, Pratap A, Agarwal S, Meyarivan T: **A fast and elitist multiobjective genetic algorithm: NSGA-II.** *IEEE Trans Evol Comput* 2002, **6**:182–197, <https://doi.org/10.1109/4235.996017>.
- Blaschke T, Engkvist O, Bajorath J, Chen H: **Memory-assisted reinforcement learning for diverse molecular de novo design.** *J Cheminf* 2020, **12**:68. <https://doi.org/10.1186/s13321-020-00473-0>.
- Yuan Y, Pei J, Lai L: **LigBuilder V3: a multi-target de novo drug design approach.** *Front Chem* 2020, **8**, <https://doi.org/10.3389/fchem.2020.00142>.
- Billsland AE, McAulay K, West R, Pugliese A, Bower J: **Automated generation of novel fragments using screening data, a dual smiles autoencoder, transfer learning and syntax correction.** *J Chem Inf Model* 2021, **61**:2547–2559, <https://doi.org/10.1021/acs.jcim.0c01226>.
- Liu X, Li P, Meng F, Zhou H, Zhong H, Zhou J, Mou L, Song S: **Simulated annealing for optimization of graphs and sequences.** *Neurocomputing* 2021, **465**:310–324, <https://doi.org/10.1016/j.neucom.2021.09.003>.
- Cofala T, Elend L, Mirbach P, Prellberg J, Teusch T, Kramer O: **Evolutionary multi-objective design of SARS-CoV-2 Protease Inhibitor candidates.** *Parallel Problem Solving from Nature - PPSN XVI. PPSN 2020. Lecture Notes in Computer Science*, vol. 12270; 2020:357–371, https://doi.org/10.1007/978-3-030-58115-2_25.
- Verhellen J: **Graph-based molecular Pareto optimisation.** *Chem Sci* 2022, **13**:7526–7535, <https://doi.org/10.1039/D2SC00821A>.
- The first application of NSGA-III in multi-objective compound design and making use of the dominated hypervolume as a novel measure of the effectiveness in multi-objective optimization.
- Devi RV, Siva Sathya S, Coumar MS: **Multi-objective biofilm algorithm (MOBifi) for de novo drug design with special focus to anti-diabetic drugs.** *Appl Soft Comput* 2020, **96**, 106655, <https://doi.org/10.1016/j.asoc.2020.106655>.
- Blaschke T, Arús-Pous J, Chen H, Margreitter C, Tyrchan C, Engkvist O, Papadopoulos K, Patronov A: **Reinvent 2.0: an AI tool for de novo drug design.** *J Chem Inf Model* 2020, **60**: 5918–5922, <https://doi.org/10.1021/acs.jcim.0c00915>.
- Fialková V, Zhao J, Papadopoulos K, Engkvist O, Bjerrum EJ, Kogej T, Patronov A: **LibINVENT: reaction-based generative scaffold decoration for in silico library design.** *J Chem Inf Model* 2022, **62**:2046–2063, <https://doi.org/10.1021/acs.jcim.1c00469>.

22. Urbina F, Lowden CT, Culberson JC, Ekins S: **MegaSyn: Integrating generative molecular design, automated analog designer, and synthetic viability prediction.** *ACS Omega* 2022, 7:18699–18713, <https://doi.org/10.1021/acsomega.2c01404>.
23. Bung N, Krishnan SR, Roy A: **An in silico explainable multi-parameter optimization approach for de novo drug design against proteins from the central nervous system.** *J Chem Inf Model* 2022, 62:2685–2695, <https://doi.org/10.1021/acs.jcim.2c00462>.
24. Perron Q, Mirguez O, Tajmouati H, Skiredj A, Rojas A, Gohier A, Ducrot P, Bourguignon M-P, Sansilvestri-Morel P, Do Huu N, Gellibert F, Gaston-Mathé Y: **Deep generative models for ligand-based de novo design applied to multi-parametric optimization.** *J Comput Chem* 2022, 43:692–703, <https://doi.org/10.1002/jcc.26826>.
25. Guo J, Fialková V, Arango JD, Margreitter C, Janet JP, Papadopoulos K, Engkvist O, Patronov A: **Improving de novo molecular design with curriculum learning.** *Nat Mach Intell* 2022, 4:555–563, <https://doi.org/10.1038/s42256-022-00494-4>.
26. Goel M, Raghunathan S, Laghuvarapu S, Priyakumar UD: **MoleGuLAR: molecule generation using reinforcement learning with alternating rewards.** *J Chem Inf Model* 2021, 61:5815–5826, <https://doi.org/10.1021/acs.jcim.1c01341>.
MoleGuLAR combines an RNN model with an RL strategy applying multi-objective optimization through alternating the reward to optimize one property at a time.
27. Liu X, Ye K, van Vlijmen HWT, Emmerich MTM, Ijzerman AP, van Westen GJP: **DrugEx v2: de novo design of drug molecules by Pareto-based multi-objective reinforcement learning in polypharmacology.** *J Cheminf* 2021, 13:85, <https://doi.org/10.1186/s13321-021-00561-9>.
28. Bolcato G, Heid E, Boström J: **On the value of using 3D shape and electrostatic similarities in deep generative methods.** *J Chem Inf Model* 2022, 62:1388–1398, <https://doi.org/10.1021/acs.jcim.1c01535>.
A multi-objective RNN RL framework with an NSGA-II based optimization scheme using Tanimoto distance to stimulate diverse solutions
29. Pereira T, Abbasi M, Oliveira JL, Ribeiro B, Arrais J: **Optimizing blood–brain barrier permeation through deep reinforcement learning for de novo drug design.** *Bioinformatics* 2021, 37(Supplement 1):i84–i92, <https://doi.org/10.1093/bioinformatics/btab301>.
30. Khemchandani Y, O'Hagan S, Samanta S, Swainston N, Roberts TJ, Bollegala D, Kell DB: **DeepGraphMolGen, a multi-objective, computational strategy for generating molecules with desirable properties: a graph convolution and reinforcement learning approach.** *J Cheminf* 2020, 12:53, <https://doi.org/10.1186/s13321-020-00454-3>.
31. Reinforced molecular optimization with Neighborhood-Controlled grammars, in: *Advances in neural information processing systems*, Vol. vol. 33, Curran Associates, Inc.
32. Lu F, Li M, Min X, Li C, Zeng X: **De novo generation of dual-target ligands using adversarial training and reinforcement learning.** *Briefings Bioinf* 2021, 22, <https://doi.org/10.1093/bib/bbab333>.
33. Abbasi M, Santos BP, Pereira TC, Sofia R, Monteiro NRC, Simões CJV, Brito R, Ribeiro B, Oliveira JL, Arrais JP: **Designing optimized drug candidates with generative adversarial network.** *J Cheminf* 2022, 14:40, <https://doi.org/10.1186/s13321-022-00623-6>.
34. Born J, Manica M, Cadow J, Markert G, Mill NA, Filipavicius M, Janakarajan N, Cardinale A, Laino T, Martónez MR: **Data-driven molecular design for discovery and synthesis of novel ligands: a case study on SARS-CoV-2.** *Mach Learn: Sci Technol* 2021, 2, 025024, <https://doi.org/10.1088/2632-2153/abe808>.
35. Wang J, Hsieh K, Wang M, Wang X, Wu Z, de Jun J, Liao B, Zhang X, Yang B, He Q, Cao D-S, Chen X, Hou T: **Multi-constraint molecular generation based on conditional transformer, knowledge distillation and reinforcement learning.** *Nat Mach Intell* 2021, 3:914–922, <https://doi.org/10.1038/s42256-021-00403-1>.
36. Jin W, Barzilay DR, Jaakkola T: **Multi-objective molecule generation using interpretable substructures.** In *Proceedings of the 37th International conference on machine learning*. PMLR; 2020:4849–4859.
A unique approach to multi-objective generation through identifying interpretable substructures for each objective and combining them using a VAE to fulfill all objectives
37. Langevin M, Minoux H, Levesque M, Bianciotto M: **Scaffold-constrained molecular generation.** *J Chem Inf Model* 2020, 60:5637–5646, <https://doi.org/10.1021/acs.jcim.0c01015>.
38. Alberga D, Gambacorta N, Trisciuzzi D, Ciriaco F, Amoroso N, Nicolotti O: **De novo drug design of targeted chemical Libraries based on artificial intelligence and pair-based multi-objective optimization.** *J Chem Inf Model* 2020, 60:4582–4593, <https://doi.org/10.1021/acs.jcim.0c00517>.
39. Kotsias P-C, Arús-Pous J, Chen H, Engkvist O, Tyrchan C, Bjerrum EJ: **Direct steering of de novo molecular generation with descriptor conditional recurrent neural networks.** *Nat Mach Intell* 2020, 2:254–265, <https://doi.org/10.1038/s42256-020-0174-5>.
A simple conditional learning approach was proposed by constructing a cRNN. It differs from cVAEs and cTrans, as the generation is only conditioned with molecular properties without molecular encoding during training or generation.
40. Lee M, Min K: **MGCVAE: multi-objective inverse design via molecular graph conditional variational autoencoder.** *J Chem Inf Model* 2022, 62:2943–2950, <https://doi.org/10.1021/acs.jcim.2c00487>.
41. Kim H, Na J, Lee WB: **Generative chemical transformer: neural machine learning of molecular geometric structures from chemical language via attention.** *J Chem Inf Model* 2021, 61:5804–5814, <https://doi.org/10.1021/acs.jcim.1c01289>.
42. Shin B, Park S, Bak J, Ho JC: **Controlled molecule generator for optimizing multiple chemical properties.** In *Proceedings of the conference on Health, Inference, and learning*. New York, NY, USA: CHIL '21, Association for Computing Machinery; 2021: 146–153, <https://doi.org/10.1145/3450439.3451879>.
43. Bagal V, Aggarwal R, Vinod PK, Priyakumar UD: **MolGPT: molecular generation using a transformer-decoder model.** *J Chem Inf Model* 2022, 62:2064–2076, <https://doi.org/10.1021/acs.jcim.1c00600>.
44. He J, You H, Sandström E, Nittinger E, Bjerrum EJ, Tyrchan C, Czechtizky W, Engkvist O: **Molecular optimization by capturing chemist's intuition using deep neural networks.** *J Cheminf* 2021, 13:26, <https://doi.org/10.1186/s13321-021-00497-0>.
45. Gao K, Nguyen DD, Tu M, Wei G-W: **Generative network complex for the automated generation of drug-like molecules.** *J Chem Inf Model* 2020, 60:5682–5698, <https://doi.org/10.1021/acs.jcim.0c00599>.
A Generative Network Complex proposed a novel approach to generate new compounds by multi-objective optimization via gradient descent in the latent space of an autoencoder.
46. Nigam A, Pollice R, Krenn M, Gomes GdP, Aspuru-Guzik A: **Beyond generative models: superfast traversal, optimization, novelty, exploration and discovery (STONED) algorithm for molecules using SELFIES.** *Chem Sci* 2021, 12:7079–7090, <https://doi.org/10.1039/D1SC00231G>.
STONED proposes an original approach to tackle multi-objective optimization in drug design by selecting a high-scoring molecule per objective and finding median compounds close to the input molecules.
47. Yasonik J: **Multiobjective de novo drug design with recurrent neural networks and nondominated sorting.** *J Cheminf* 2020, 12:14, <https://doi.org/10.1186/s13321-020-00419-6>.
48. Iovanac NC, MacKnight R, Savoie BM: **Actively searching: Inverse design of novel molecules with simultaneously optimized properties.** *J Phys Chem* 2022, 126:333–340, <https://doi.org/10.1021/acs.jpca.1c08191>.
A VAE framework that incorporates the concepts of conditional and active learning for multi-objective optimization, where iteratively the best output of the generator is used to continue training the model
49. He J, Nittinger E, Tyrchan C, Czechtizky W, Patronov A, Bjerrum EJ, Engkvist O: **Transformer neural network-based molecular optimization using general transformations.** *J Cheminform* 2022, <https://doi.org/10.1186/s13321-022-00599-3>.

50. Elend L, Jacobsen L, Cofala T, Prellberg J, Teusch T, Kramer O, Solov'yov IA: **Design of SARS-CoV-2 main protease inhibitors using artificial intelligence and molecular dynamic simulations.** *Molecules* 2022, **27**, <https://doi.org/10.3390/molecules27134020>. 4020, number: 13 Publisher: Multidisciplinary Digital Publishing Institute.
51. Horwood J, Noutahi E: **Molecular design in synthetically accessible chemical space via deep reinforcement learning.** *ACS Omega* 2020, **5**:32984–32994, <https://doi.org/10.1021/acsomega.0c04153>.
52. C. von Lücken, C. Brizuela, B. Barán, An overview on evolutionary algorithms for many-objective optimization problems, *WIREs Data Mining and Knowledge Discovery* 9 e1267. doi: <https://doi.org/10.1002/widm.1267>.
53. Brown N, Fiscato M, Segler MH, Vaucher AC: **GuacaMol: benchmarking models for de novo molecular design.** *J Chem Inf Model* 2019, **59**:1096–1108, <https://doi.org/10.1021/acs.jcim.8b00839>. American Chemical Society.
54. Polykovskiy D, Zhebrak A, Sanchez-Lengeling B, Golovanov S, Tatanov O, Belyaev S, Kurbanov R, Artamonov A, Aladinskiy V, Veselov M, Kadurin A, Johansson S, Chen H, Nikolenko S, Aspuru-Guzik A, Zhavoronkov A: **Molecular sets (MOSES): a benchmarking platform for molecular generation models.** *Front Pharmacol* 2020, **11**, <https://doi.org/10.3389/fphar.2020.565644>.
55. Reeves S, DiFrancesco B, Shahani V, MacKinnon S, Windemuth A, Brereton AE: **Assessing methods and obstacles in chemical space exploration.** *Applied AI Letters* 2020, **1**:e17, <https://doi.org/10.1002/ail2.17>.
56. Schwaller P, Vaucher AC, Laplaza R, Bunne C, Krause A, Corminboeuf C, Laino T: **Machine intelligence for chemical reaction space.** In *WIREs Computational Molecular Science*; 2022, e1604, <https://doi.org/10.1002/wcms.1604>.
57. Laato S, Tiainen M, Islam AKMN, Mäntymäki M: **How to explain ai systems to end users: a systematic literature review and research agenda.** *Internet Res* 2022:1–31, <https://doi.org/10.1108/INTR-08-2021-0600>.
58. Emmerich M, Yang K, Deutz A, Wang H, Fonseca CM: **A multicriteria generalization of bayesian global optimization.** In *Advances in Stochastic and Deterministic Global optimization.* Springer; 2016:229–242.
59. Mervin LH, Johansson S, Semenova E, Giblin KA, Engkvist O: **Uncertainty quantification in drug design.** *Drug Discov Today* 2021, **26**:474–489, <https://doi.org/10.1016/j.drudis.2020.11.027>.
60. Nigam A, Pollice R, Hurley MFD, Hickman RJ, Aldeghi M, Yoshikawa N, Chithrananda S, Voelz VA, Aspuru-Guzik A: **Assigning confidence to molecular property prediction.** *Expet Opin Drug Discov Sep.* 2021, **16**:1009–1023, <https://doi.org/10.1080/17460441.2021.1925247>.