Computers in Biology and Medicine 145 (2022) 105403

Contents lists available at ScienceDirect

Computers in Biology and Medicine

journal homepage: www.elsevier.com/locate/compbiomed

Generative machine learning for de novo drug discovery: A systematic review

Dominic D. Martinelli

Independent researcher, United States of America

ARTICLE INFO

Keywords:
Artificial intelligence
Cheminformatics
Generative machine learning
de novo drug discovery

ABSTRACT

Recent research on artificial intelligence indicates that machine learning algorithms can auto-generate novel drug-like molecules. Generative models have revolutionized de novo drug discovery, rendering the explorative process more efficient. Several model frameworks and input formats have been proposed to enhance the performance of intelligent algorithms in generative molecular design. In this systematic literature review of experimental articles and reviews over the last five years, machine learning models, challenges associated with computational molecule design along with proposed solutions, and molecular encoding methods are discussed. A query-based search of the PubMed, ScienceDirect, Springer, Wiley Online Library, arXiv, MDPI, bioRxiv, and IEEE Xplore databases yielded 87 studies. Twelve additional studies were identified via citation searching. Of the articles in which machine learning was implemented, six prominent algorithms were identified: long short-term memory recurrent neural networks (LSTM-RNNs), variational autoencoders (VAEs), generative adversarial networks (GANs), adversarial autoencoders (AAEs), evolutionary algorithms, and gated recurrent unit (GRU-RNNs). Furthermore, eight central challenges were designated: homogeneity of generated molecular libraries, deficient synthesizability, limited assay data, model interpretability, incapacity for multi-property optimization, incomparability, restricted molecule size, and uncertainty in model evaluation. Molecules were encoded either as strings, which were occasionally augmented using randomization, as 2D graphs, or as 3D graphs. Statistical analysis and visualization are performed to illustrate how approaches to machine learning in de novo drug design have evolved over the past five years. Finally, future opportunities and reservations are discussed.

1. Introduction

Machine learning, a subtype of artificial intelligence constituting an algorithm able to improve itself independent of human intervention, has gained immense popularity in the medical industry since its advent [1]. The technique can automate processes that would otherwise consume a considerable amount of time and resources, making accessible, efficient healthcare a more realistic objective. In the field of medicinal chemistry, computational methods have been applied in many areas of the drug discovery process, including the evaluation of compound similarity, molecule classification, and bioactivity prediction [2]. Further attempts to yield safe and efficacious drug-like compounds indirectly, for example, by identifying propitious targets rather than concentrating on the properties of small molecules, have been explored through the integration of machine learning. Numerous studies have enabled the classification of promising macromolecular targets, even aiding in the discovery of drugs now approved by the FDA [3,4]. Virtual screening, however, can only contribute so much. Because virtual screening requires the navigation of large chemical spaces, exploring them can be computationally expensive [5,6]. If the costs incurred by such methods exceed their benefits, it cannot be responsibly argued that computational approaches are inherently efficient.

Recently, more advanced machine learning models devised to generate new information have been proposed, among them recurrent neural networks (RNNs) [7], generative adversarial networks (GANs) [8], and variational autoencoders (VAEs) [9]. These techniques have been related to a wide range of problems, including sentimental text generation, music composition, inpainting, and trajectory prediction, among many others [10–13]. With an expanding body of evidence to support the application of machine learning to complex problems, the potential for mechanistic, de novo compound generation has gained extensive attention. Various models have since been devised to accomplish this task, each offering insight into the relative value of different generative methods. In the context of drug discovery, machine learning-informed drug development constitutes an improvement to manual virtual screening in that it is both automated and adaptable

E-mail address: mail.dmartinelli@gmail.com.

https://doi.org/10.1016/j.compbiomed.2022.105403
Received 14 January 2022; Received in revised form 10 March 2022; Accepted 11 March 2022
Available online 13 March 2022

without imposing a significant computational burden [14]. As such, these generative models could revolutionize the future of pharmaceutical engineering. Generating novel drug-like substances in silico is a multifaceted issue that cannot be addressed without an understanding of the great volume of research currently available. Evidently, the interconnectivity between medical research and informatics has substantial implications for the expanding efficiency and precision of pharmacological innovation, introducing opportunities to extend novel treatments at a reduced cost. Identifying areas in need of improvement, along with existing solutions, is essential to the attainment of this objective. Models capable of generating new drug-like compounds demonstrate great promise, as indicated by the evolving library of studies dedicated to assessing their performances. Moreover, the ability to emulate the complexities of chemical space using computational tools would denote that similarly complicated questions in medicine are within the scope of intelligent algorithms.

In this systematic review, the evolution of the use of generative machine learning algorithms in de novo drug discovery over the past five years is examined comprehensively. Answers to the following research questions were pursued:

1. Which machine learning models are implemented to generate novel molecular structures de novo?
2. Which challenges related to this machine learning problem have researchers addressed?
3. How are molecular structures encoded in the context of machine learning?

While some recent reviews have addressed the role of artificial intelligence in drug development, these broad overviews do not provide detailed insight regarding generative methods, instead limiting discussion on studies of generative drug design to brief summaries of selected developments [15,16]. The breadth of models discussed is too narrow to reflect the full scope of this machine learning task, especially considering the rate at which novel solutions have emerged [17]. Here, a diverse collection of studies, representative of recent advancements in the discipline, is presented in an exhaustive systematic review. Notably, changes in the use of specific machine learning models, molecular encoding techniques, and challenges addressed over the past five years are described by means of statistical analysis. This technique facilitates the presentation of the appreciable volume of information associated with this application of machine learning. Developers and researchers may use this publication to inform their methodologies and approaches to de novo drug discovery in silico. Additionally, inferences may be made about the evolving relationship between medicine and artificial intelligence.

2. Conceptual background

To understand the answers to the proposed review questions, a summary of the terminology applied in computational chemistry and generative machine learning is provided.

2.1. Generative machine learning models

2.1.1. RNN

In 1986, David Rumelhart introduced neural backpropagation, a learning procedure in which weights are adjusted based on the calculated loss function gradient such that the output values might better resemble the input [7]. This principle would inspire the RNN architecture (Fig. 1a.). RNNs possess a memory, or internal state, which enables them to detect patterns in inputs of variable lengths. RNNs may also have additional storage states controlled by a neural network, either internally or externally. LSTM and gated GRUs are two common examples of these controlled storage states. In LSTM-RNN (Fig. 1b.), the vanishing gradient problem often encountered in backpropagation is less likely to occur. Similar to LSTM, the GRU model differs only in that it incorporates fewer parameters and excludes an output gate (Fig. 1c.).

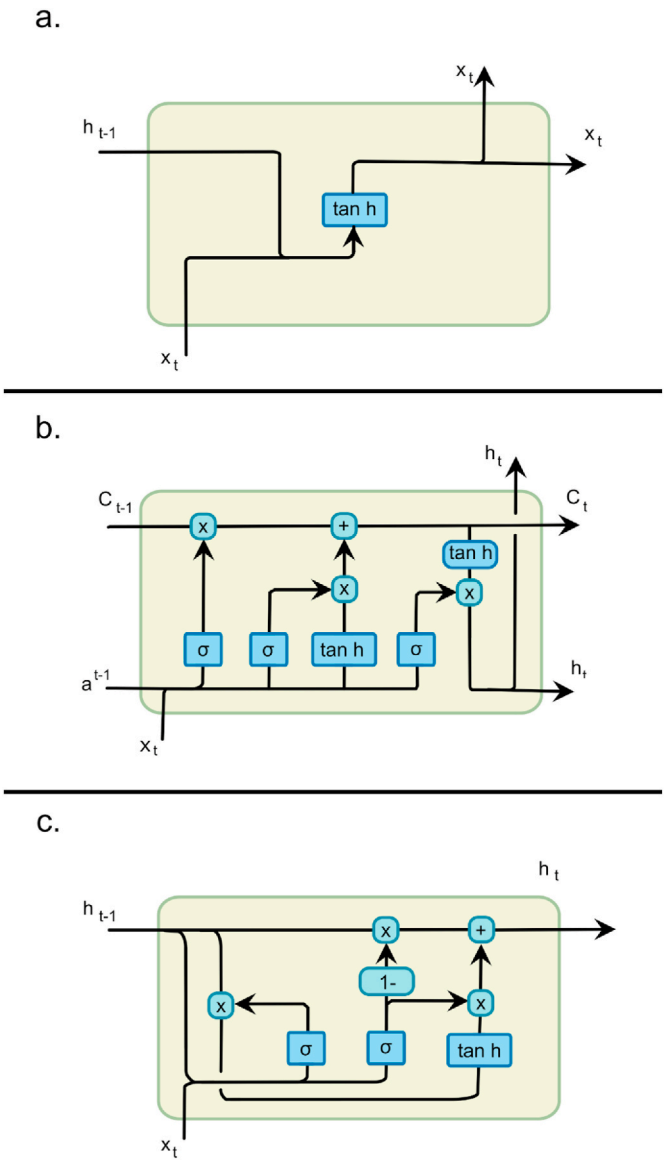

Fig. 1. Three varieties of recurrent neural networks, including (a) a standard RNN, (b) an LSTM-RNN, and (c) a GRU-RNN. The inputs are represented by x_t, the current input; C_{t-1}, the cell state; and h_{t-1}, the hidden state. The outputs are C_t, the next cell state, and h_t, the new output. The nonlinearities are represented by σ and tan h, the sigmoid and hyperbolic tangent layers, respectively. Finally, the × and + operators indicate pointwise multiplication and addition.

2.1.2. GAN

The GAN machine learning architecture was proposed in June 2014 by Ian Goodfellow and his colleagues [8] (Fig. 2). In a GAN, two neural networks compete in a zero-sum game. In other words, when one of the algorithms succeeds at the specified task, or "wins," the other "loses." The first of the two competitors, called a generative network, attempts to mimic the input, while the discriminative network determines whether the data was generated artificially. The objective of the generative network is to "fool" the discriminator such that the synthesized images are classified as components of the true data distribution. When the composition proposed by the generator is incorrectly classified as genuine by the discriminator, it is considered to be at an advantage, while the opposite is true for the discriminator network. "Mode collapse" is a common problem in GANs, in which they fail to detect

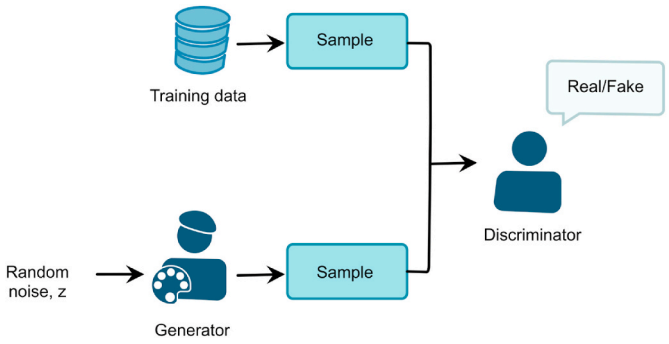

Fig. 2. A generative adversarial network. The generator assembles the random noise such that it, ideally, resembles the training data. The discriminator then categorizes each sample.

several modes from the input data. A GAN suffering mode collapse will only generate a very limited range of outputs, if not merely a single output.

2.1.3. AE

Autoencoders (AEs) are used to accomplish unsupervised machine learning tasks. While AEs in general have a wide variety of applications in machine learning, variational AEs (VAEs) are particularly conventional generative algorithms (Fig. 3). Introduced by Diederik P. Kingma and Max Welling, the VAE is able to encode and decode information just as a standard autoencoder would [9]. However, unlike the simple autoencoder, it is a probabilistic model, rather than a deterministic model. Input is encoded as a latent distribution. Before decoding the input, the latent representation is sampled from the distribution. The output depicts the decoded manifestation of the sampled latent representation.

Adversarial AEs (AAEs) are another form of probabilistic AE that use a GAN to match the aggregated posterior of the latent representation to the prior distribution [18] (Fig. 4). The original AAE paper describes a model that includes both a traditional reconstruction error criterion and an adversarial training criterion. A sample is taken from a distribution chosen by the user and evaluated by the discriminator for its resemblance to the training data.

2.1.4. Evolutionary algorithm

Evolutionary algorithms are population-based models, often inspired by biological evolution, constructed to solve optimization problems [19]. Several operators (Fig. 5) can be applied to individual data points to modify their characteristics and evaluate their fitness for a given objective. Genetic algorithms are a subtype of evolutionary algorithms, in which solutions to an optimization problem are encoded as strings and subjected to operators. Another application of the evolutionary algorithm is genetic programming, in which a program's fitness is assessed in terms of its ability to complete a computational task.

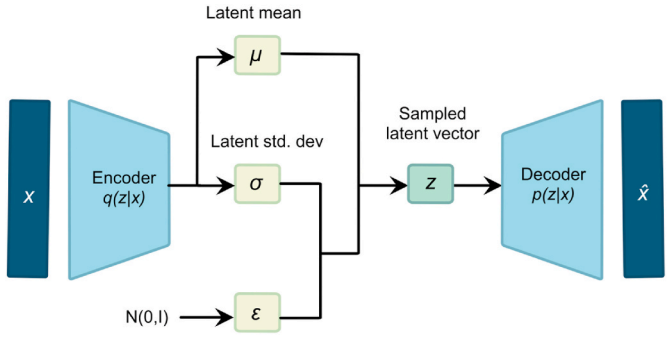

Fig. 3. A variational autoencoder. As indicated, the encoder and decoder are probabilistic, not deterministic. The sampled latent vector is a compressed version of the input. N(0,1) represents the normal distribution. The error term is denoted as ε.

2.2. Molecular encoding

2.2.1. String-based representations

The Simplified Molecular Input Line Entry System (SMILES) was established by David Weninger as a means of representing individual molecules as strings [20]. The system, exemplified in Fig. 6, was inspired by molecular graph theory. SMILES provide information about a chemical's atoms, bonds, branches, and cyclic structures. A canonical SMILES string does not provide information about isotopism and stereochemistry, while isomeric SMILES strings do. Additionally, aromaticity may be expressed by altering the case of an atom. Although each SMILES string only corresponds to one molecule, a unique molecule may possess several SMILES representations. Other authors have elaborated on Weninger's SMILES syntax. One notable derivative of the original SMILES system, SELFIES, was created to mitigate the problem of random SMILES invalidity [21]. SELFIES strings may be more suitable for processing by machine learning models, especially generative algorithms, as they are completely robust. This means that even randomly generated SELFIES strings will depict valid molecules.

Binary molecular fingerprints may also represent molecules, especially when attempting to compare chemical structures objectively. The Molecular Access System (MACCS) key fingerprint is a 166-key long string of binary values [22]. MACCS keys can depict the features and structural characteristics of individual molecules without incurring a high computational burden or excessive complexity. Many of the features characterized, like aromatic ring count, are properties of interest in a pharmacological context [23,24].

2.2.2. 2D graphs

2D molecular graphs consist of a set of nodes V and a set of edges E, where V represents the set of atoms contained in a molecule, and E represents the set of bonds linking those atoms together [25]. An adjacency matrix (Fig. 7a) may indicate the position of a given atom in a compound. Bonds are depicted by an edge features matrix (Fig. 7b), while atoms are represented by a node features matrix (Fig. 7c). Atoms, bonds, and adjacencies are designated as one-hot encoded matrices, but integer encoded matrices may be used with less complex molecules as well.

2.2.3. 3D graphs

Several variations of 3D molecular graphs have been proposed for use in machine learning algorithms, but their fundamental elements are nevertheless inspired by molecular graph theory. Bond angles, bond lengths, and dihedral angles may be described as diagonal matrices [26]. Other models which incorporate feature, adjacency, and relative position matrices incorporate existing knowledge about the quantitative basis for bond polarity, electronegativity, and 3D conformation to construct graphs suitable for graph convolutional networks [27]. Alternatively, the coordinates of each atom in the chemical space can be defined explicitly [28].

2.3. In silico validation

While comparative benchmarks designed specifically for generative models will be discussed in detail, the principles upon which these systems are contrived must be reviewed. Databases curated to address target-ligand interactions enable the training of algorithms according to a defined set of attributes. Compound-focused libraries, such as ZINC and DrugBank, specify the clinical applicability of approved drugs and drug candidates [29,30]. Predicting molecule activity congruent to its structure may be accomplished by studying these collections of drugs, rather than exploring the entire chemical space, improving efficiency

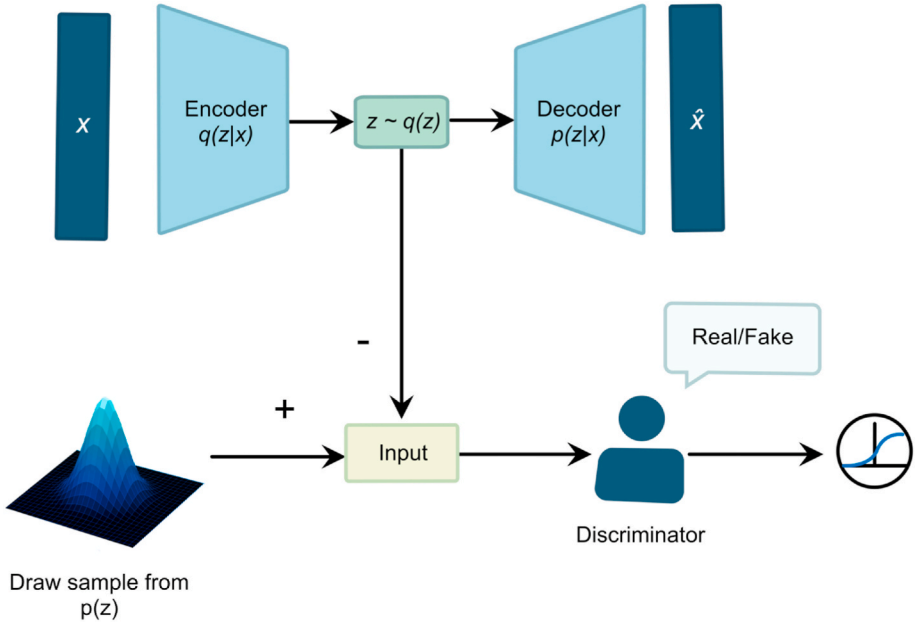

Fig. 4. An adversarial autoencoder. The top row is a basic autoencoder, while the bottom row represents the adversarial piece of the model in which the discriminator determines whether the generated output resembles the training data.

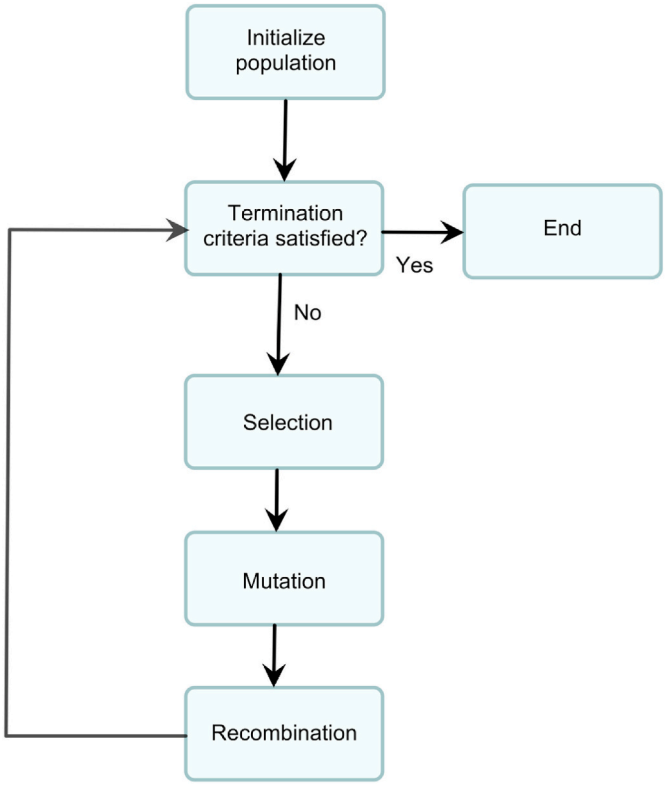

Fig. 5. The basic steps executed in an evolutionary algorithm. Termination criteria could be defined in many ways, such as optimization or a set maximum number of iterations.

and specificity. Other databases, such as the IUPHAR Guide to Pharmacology and ChEMBL libraries, offer greater detail on the dynamic relationships between targets and ligands [31,32]. They include profiles curated for individual targets, which is convenient in assessing potential interactions between drugs and the role of protein function in ligand binding. Furthermore, target-focused libraries like the Transporter

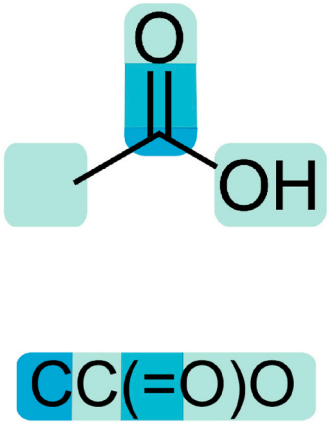

Fig. 6. Acetic acid and its corresponding SMILES representation. The double bond is portrayed by an equal sign and the branch structure is represented by parentheses.

Classification Database (TCDB) and the Variability of Drug Transporter Database (VARIDT) may inform a machine learning algorithm of the function of structural characteristics of proteins in drug efficacy [33,34]. However, as de novo compound design concerns small molecules, data about the target alone does not provide the chemical information necessary to synthesize original ligands. Therefore, it is vital to use these databases as complementary agents.

Generating novel molecular structures is not a particularly difficult task in and of itself; the challenge lies in synthesizing compounds that are both structurally plausible and drug-like. If a compound is not chemically possible, it cannot be evaluated, regardless of how well it seems to perform in silico. Multi-property optimization, as defined in this study, refers to the capacity of a model to satisfy these various conditions. Similarity, which is often determined through the use of molecular fingerprints and pharmacophoric descriptors, is a common metric used to determine validity [35]. When receptor features cannot be encoded, comparing generated molecules to compounds with known bioactivity with respect to a target enables the prediction of biophysical mechanisms without specifying information about protein structure

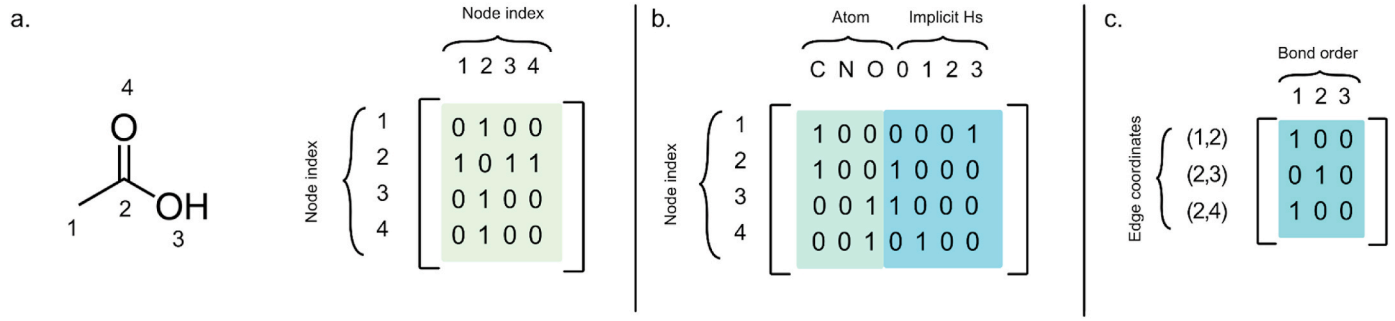

Fig. 7. Acetic acid encoded with (a) an adjacency matrix, (b) a node features matrix with information about the atom types and number of implicit hydrogen atoms, and (c) an edge features matrix representing the bond orders.

[36]. Broadly, the qualities of interest assessed in multi-property models include absorption, distribution, metabolism, excretion, and toxicity (ADMET). Depending on the training data and the specified bioactivity, several measures of therapeutic potential can be utilized. Binding affinity, correspondence to Lipinski's rule of five, and solubility are among the most frequent measures of ADMET satisfaction [37,38]. In many of the established databases used to train generative models, like ChEMBL, entries indicate both molecular structure and ADMET descriptors. Therefore, whether a system adequately optimizes multiple characteristics depends on the fulfillment of these requirements.

Docking and molecular dynamics simulations (Fig. 8) are studied to evaluate interactions between ligands and their environments or macromolecular targets. In molecular docking, a ligand is positioned relative to a specified target and scored as a function of its affinity for the binding site. Metrics provided to compare binding affinity are provided in numerous docking software interfaces. Molecular docking, as a form of compound validation, can assist in the accurate prediction of chemical behavior in an experimental setting [39,40]. Molecular dynamics simulations may elucidate binding mechanisms and advise the selection of more selective ligands [41–43]. These simulations are also conducted with the objective of making informed predictions about the expected behavior of drug candidates when acting on a target.

3. Systematic literature review

To ensure a standardized review procedure, the Preferred Reporting Items for Systematic reviews and Meta-analyses (PRISMA) 2020 statement informed the methodology used in this review [44]. The statement includes guidelines for authors, specifically discussing the measures that should be taken to ensure complete transparency in conducting systematic reviews. All items stated will therefore be addressed prior to formal analysis.

3.1. Eligibility criteria

Because machine learning methods contrived to generate novel molecules automatically have only recently expanded, a five-year time period was deemed appropriate for the present review. For example, the de novo drug design software Design of Genuine Structures (DOGS), introduced in 2012, does not utilize any generative machine learning technique, although it would ultimately inspire subsequent experiments [45]. Until 2017, the explicit application of generative molecular models to de novo pharmacology was not stated. Notably, the publications of Gómez-Bombarelli et al. and Segler et al. were the first observed to reference each other concerning this problem [46,47]. Considering the relative novelty of this field, the selected time period was chosen to accommodate the sparsity of relevant content prior to 2017.

Only articles found to satisfy the following criteria were included in the review:

1. The study is written in English.
2. The study was published between January 2017 and January 2022.
3. The study is a full-text research article or review.
4. The study concerns the use of machine learning to generate small, drug-like molecules de novo.

3.2. Information sources

The following online databases were searched for eligible articles: (1) PubMed, an archive of biomedical literature maintained by the U.S. National Institute of Health; (2) ScienceDirect, a peer-reviewed article database containing full-text articles across a wide range of disciplines published in Elsevier journals; (3) Springer, an international source of scientific documents from journals, books, proceedings, and protocols; (4) Wiley Online Library, a multidisciplinary source of journal

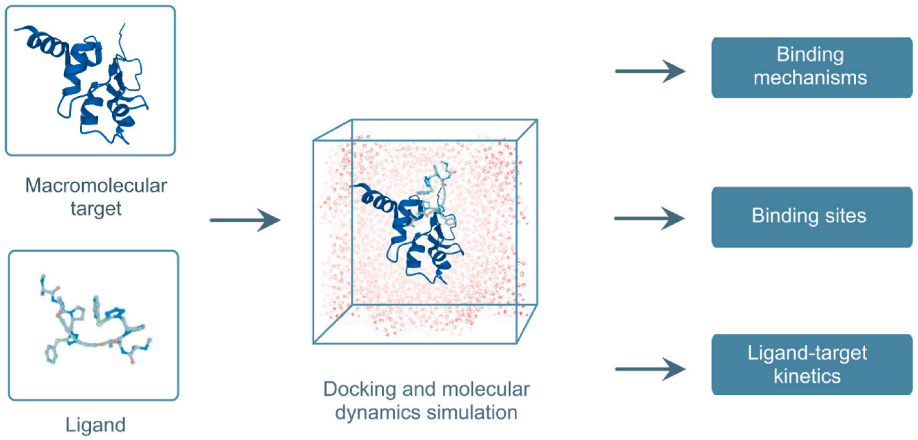

Fig. 8. A schematic representation of the rationale behind molecular dynamics simulations as a means of validating ligand-target interactions.

publications, books, and other research documents; (5) arXiv, an open-access database of scholarly articles; (6) MDPI, an open-access publisher of peer-reviewed studies; (7) bioRxiv, a distribution service for article preprints; and (8) IEEE Xplore, a digital library maintained by the Institute of Electrical and Electronics Engineers containing over five million scientific and technical documents. All databases provide domain-relevant information.

3.3. Search strategy

In each of the aforementioned databases, an Advanced Search filter was applied to limit search results to articles published since January 2017. To ensure complete transparency of information, only full-text articles were screened manually. The search queries were written as Boolean statements as follows:

1. "Machine learning" AND "de novo" AND "drug";
2. "Machine learning" AND "de novo" AND "molecul*".

3.4. Study selection

After searching the specified databases using the filters described in section 2.3, a total of 1,402 articles were collected. To assess whether the inclusion criteria mentioned were met, the abstracts and titles of each article were reviewed manually. If the abstract and title of an article suggested that the document could be within the scope of the present systematic literature review, the text was read in its entirety. Studies irrelevant to the research questions were discarded from consideration. 1,301 studies were deemed ineligible for the present systematic literature review. Fourteen of the remaining 101 studies were duplicate articles. Therefore, the initial search of the literature yielded 87 full-length articles. Twelve additional publications were included in the review on account of their frequent citations by the other texts identified. In total, 96 studies were found eligible for this review. The PRISMA flow chart corresponding to the present methodology is shown in Fig. 9.

3.5. Data collection

In accordance with the research questions proposed during this systematic review, the data items shown in Table 1 were collected for each article selected. Results were organized in an Excel spreadsheet.

Table 1
The data recorded for each individual document in an Excel spreadsheet.

Item	Description
Title	Title of the document
Publication Year	Year of publication
Model(s)	Machine learning model(s) used or designed
Molecular Representation	Molecular encoding format
Challenges	Challenges addressed, independent of whether a solution was proposed

Titles were collected to facilitate the identification of each study. The publication year of each article was also included to enable comparative analysis of the distribution of other article characteristics over the past five years. For original research articles, the model(s) used were recorded to evaluate trends and improvements to machine learning models used in generative drug discovery. Because generative models depend so heavily on the format of input data, the encoded representation utilized to format each molecule was recorded as well. For every article in which difficult aspects of in silico de novo drug design were disclosed, the specific concerns raised were recorded as challenges. Note that, likewise, these challenges imply proposed countermeasures to express difficulties. For review articles, the models and molecular representations discussed were not included in the final statistical analysis but were still listed in their respective columns to observe which among them were most frequently discussed throughout all the articles.

3.6. Study quality assessment

To assess the quality of each article selected, the Critical Appraisal Skills Programme (CASP) Qualitative Studies and Systematic Review Checklists were modified [48,49]. The modified CASP Qualitative Studies Checklist questions (Table 2) were used to evaluate experimental research articles written either to introduce new models or to make use of existing models such that their performances might be analyzed. Questions in the original CASP Qualitative Studies Checklist referencing human and animal participants were removed, as no such procedure was relevant to the articles in question. Review or commentary articles were critiqued with the modified CASP Systematic Review

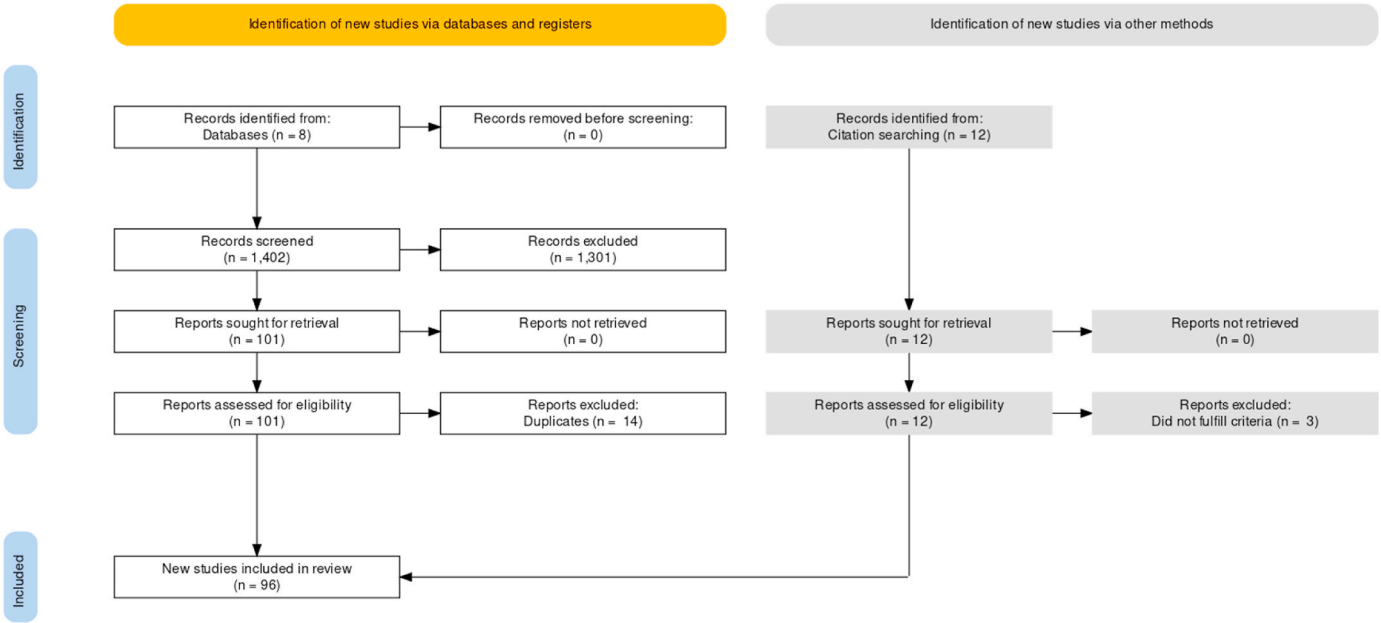

Fig. 9. The PRISMA flow diagram demonstrating the filtering of documents.

Table 2
The modified CASP Qualitative Studies Checklist items referenced to evaluate articles in which a generative model was utilized or evaluated.

Checklist item	Examples
Was there a clear statement of the aims of the research?	Did the authors state how they expected to improve other methods, or did they introduce a new approach? Were models executed to test a comparative metric or to examine an area in which present models could improve?
Is a qualitative methodology appropriate?	Were the qualities of generated models discussed in context? Were any practical implementations of the model demonstrated? How were generated molecules compared with existing drugs?
Was the research design appropriate to address the aims of the research?	If a new model was designed, was the computational basis for the algorithm rationalized? Were metrics clearly defined and justified?
Was the data collected in a way that addressed the research issue?	Which qualities were studied? Were they relevant to the aims of the research? Were appropriate databases used?
Was the data analysis sufficiently rigorous?	Did the authors exercise full transparency in presenting the data? Was the data deliberately manipulated such that readers might be deceived? Were the tools or calculations used to collect data stated?
Is there a clear statement of findings?	Were the data explained and applied to real-world contexts? Was the method evaluated relative to numerous tasks? Were the data represented honestly?
Were the contributions of the study addressed?	Were findings compared with those from previous studies? Did the authors address how their findings might be applied to future experiments or practical contexts?
Are the limitations of the study stated?	Did the paper expand upon areas in which the findings could be biased or limited? Were prospects for future methodological enhancements considered?
Has the researcher stated any conflicts of interest or lack thereof?	Was a transparent statement of conflict of interests, either present or absent, included in the article?

Table 3
The modified CASP Systematic Review Checklist items referenced to evaluate articles in which a review was conducted.

Checklist item	Examples
Was there a clear statement of the aims of the review?	Did the authors state how they expected to influence the practice of machine learning in drug discovery? Was the motivation for the paper explained?
Did the authors look for the right type of papers?	Were articles reviewed on-topic? Were published research articles sought?
Were relevant studies included?	Did each citation coherently relate to the aims of the review? Were findings taken out of context?
Were studies analyzed reliably sourced?	Did the reviewer verify the credibility of the source material? Did all studies provide sufficient evidence, even to support subjective claims?
If the results of the review have been combined, was it reasonable to do so?	If papers were categorized, was classification justified? Were criteria for including or excluding articles specified?
Is there a clear statement of findings?	Was the content of analyzed articles summarized honestly? Were all aims of the study addressed? Was statistical analysis performed to quantify the results?
Were the contributions of the study addressed?	Were findings compared with those from previous studies? Did the authors address how their findings might be applied to future experiments or practical contexts?
Were references cited when appropriate?	If subjective claims were made, was sufficient evidence provided to justify them? Was evidence reliable? Did the review provide a balanced perspective on the issue?
Has the researcher stated any conflicts of interest or lack thereof?	Was a transparent statement of conflict of interests, either present or absent, included in the article?

Checklist questions (Table 3). Because not all review or commentary articles were systematic reviews, many of the checklist items were irrelevant to the subject.

Almost all criteria were satisfied by every article chosen. Limitations do exist, however, in that not all articles were published without an acknowledgment of competing interests or lack thereof. Some experimental research articles failed to explicitly mention study impediments or future directions.

4. Results and discussion

4.1. Machine learning models

4.1.1. RNN

LSTM-RNN models were by far the most widespread generative architectures used in machine-informed de novo drug design. Olivecrona et al. trained an RNN employing RL to fine-tune the network, significantly improving previous efforts to generate molecular graphs with an RNN [50]. As one of the first attempts to produce a generative model to aid in drug discovery, their model is regularly compared to alternative RNN architectures. Soon after, Popova et al. advanced ReLeaSE, a stacked LSTM-RNN QSAR model, correspondingly implementing RL to optimize specified properties [51]. They justified their use of RL by citing findings that the approach could reduce bias in chemical structure generation, setting a precedent in conjunction with Olivecrona et al. Yang et al. [52] and Ertl et al. [53] demonstrated that RNNs could innovate molecular structures with an efficiency superior to that of VAEs. This characteristic is often requisite for drug discovery, especially if a given disease poses an immediate danger to those affected. Furthermore, the former team distinguished their strategy from previous methods by specifying the model's ability to suggest original structures without relying on predefined fragments. Unlike Yang et al., Ertl et al. did not perform a Monte Carlo tree search. Both studies resulted in generated drug-like molecules of high synthetic feasibility. Segler et al. [47] used an LSTM-RNN model to attempt to reproduce input chemical structures. They were able to fine tune the model's performance by training it with input molecules with known activity against an individual target and found that this manner of fine-tuning could refine generated molecules such that they will exhibit the desired bioactivity.

Gupta et al. fine-tuned an LSTM-RNN network with TL, an approach that permitted the generation of valid molecules even without enumerating a virtual compound library [54]. Like Popova et al., the team emphasized the importance of coupling the RNN with an external learning technique in order to reduce bias. Notably, their publication was the first example of fragment growing through an RNN. Merk et al. determined that an RNN could generate valid molecules independent of explicitly provided chemical rules [55]. Hence, the need to develop models encoded with domain expertise may be counterproductive and unnecessary. Their model displayed impressive generalizability and sensitivity to fine-tuning, another instance of the promise of artificial intelligence in target-specific ligand design. Bjerrum and Sattarov [56] compared the performance of LSTM-RNN models using autoencoders and heteroencoders. Following the technique proposed by Segler et al., Bruns et al. applied constructive machine learning to train an LSTM-RNN algorithm to construct molecules able to induce chemotaxis [57]. Two synthetically accessible compounds, both successful in producing the desired phenotypic effects, were discovered, proving that RNNs may be trained to generate molecules with bioactive properties beyond those unique to protein targets. DeepCOMO, introduced by Yonchev and Bajorath [58], was formulated to incorporate the analog functionality of the compound optimization monitor (COMO) method

and TL with the generative capability of an LSTM-RNN model. Whereas Gupta et al. adopted TL in conjunction with an LSTM to generate a diverse library of molecules, Yonchev and Bajorath's use of analog series enabled the simultaneous optimization and generation of novel drug-like substances.

Various authors utilized LSTM-RNN networks to generate novel kinase inhibitors [59–63]. Another practical use of an LSTM-RNN model was implemented by Santana and Silva-Jr [64], as well as by Amilpur and Bhukya [65] in a separate study, to predict inhibitors of SARS-CoV-2 protease inhibitors. These studies imply that the application of generative machine learning to individual macromolecular targets could be feasible ex silico. Santana and Silva-Jr analyzed 20 hit compounds through molecular docking, observing nine with binding positions similar to those assumed by experimental molecules. While the results obtained by Amilpur and Bhukya were similarly auspicious, the use of three complementary models—one to assemble general chemical structures, one to fine-tune chemical properties, and a final classification algorithm—by Santana and Silva-Jr may enhance prediction accuracy. REINVENT, an LSTM-RNN model, was evaluated by Thomas et al. [66] and was later fine-tuned by Blaschke and Bajorath [67]. It was concluded that alterations in the set of molecules used to train an RNN could lead to the generation of molecules exhibiting preferred characteristics similar to those of familiar molecules, as demonstrated in previous investigations of fine-tuned datasets. Moret et al. [68], using an LSTM-RNN algorithm, ascertained that external compound scoring functions were not imperative when model-intrinsic sampling was incorporated. Van Deursen et al. [69] administered a series of hidden, bidirectional RNN units to form Generative Examination Networks (GEN) as a novel training method for RNNs. Their observations suggest that, in relation to the GRU layers, LSTM layers yielded a greater proportion of valid SMILES strings. Likewise, bidirectional GRU layers performed poorly in comparison to both uni- and bidirectional LSTM layers. An LSTM-RNN was used in conjunction with a beam search algorithm by Bai et al. [70]. The algorithm conserved structural components necessary for the preservation of bioactivity while simultaneously maintaining compound diversity. Beam search sampling was also considered a favorable alternative to external compound prioritization.

The RNN-GRU architecture, while certainly less universal than LSTM-RNN, was explored as well. Zheng et al. [71] designed an RNN-GRU network informed by molecular stereochemistry and bioactivity. In the study, the productive consequences of TL integration were affirmed. Derivatives of recognized bioactive substances were more prominent in the generated library after TL was applied. The same network structure was proposed by Popova et al. [72] in their MolecularRNN, intended to compose molecules with specific properties. Unlike previous models, MolecularRNN employs a valency-based rejection sampling process, a feature that guarantees 100% structural validity. Blaschke et al. [73] harnessed memory-assisted RL to train an RNN-GRU model. Compared to standard RL algorithms, memory-assisted RL mitigated the problem of insufficient diversity in recommended molecules. Another case study completed by Liu et al. [74], in which DrugEx was employed, implicated the use of an RNN-GRU with RL and a special exploration strategy to discover ligands against the adenosine A2A receptor. This exploration method involves the coupling of two RNN models: one, deemed the "exploitation network," and the alternative "exploration network." This technique improved similarity between the artificially generated data without rewarding the production of molecules identical to those in the training set. After van Deursen et al. published a comparative analysis of GRU and LSTM layers in which the performative superiority of the LSTM architecture was substantiated, the use of GRU models declined. In spite of this, the findings of Popova et al. and Liu et al. should continue to direct the framework of LSTM-RNN models.

4.1.2. AE

VAEs were the second most prevalent class of models used to generate drug-like molecules de novo. Like RNNs, they were one of the first types of generative models to be implemented. Gómez-Bombarelli et al. [46] introduced the technique in 2016, inspiring other researchers to elaborate upon their conclusions. Because the database included in the study was limited to molecules containing no more than nine heavy atoms, their results did not represent drug design as a broader discipline. The conditional VAE introduced by Lim et al. [75] illustrates one such example. Their CVAE is distinct from the traditional VAE in that it is able to alter the encoding and decoding processes such that certain conditions will be met by embedding conditional information in the objective function of the VAE. Assouel et al. [76] manufactured DEFactor with the same intention–to facilitate conditional graph generation. Nevertheless, the model described did not resolve the issue of inferior graph size, as observed in the original work by Gómez-Bombarelli et al. GraphVAE was created by Simonovsky and Komodakis [77] to increase the efficiency of small molecular graph generation. Graph size continued to present a limitation. However, the probabilistic graph decoder successfully circumvented the problem of non-differentiability identified in the initial paper. Samanta et al. [78] proposed NeVAE, a VAE capable of generating molecular graphs adapted to specified characteristics. Compared to Bayesian optimization and established RL techniques, the output of NeVAE comprises substances with property values improved by over 121%. Li et al. [79] organized DeepScaffold in consonance with the VAE architecture. Confirming the earlier conclusions of Merk et al., DeepScaffold implicitly developed an awareness of chemical laws. Additionally, the model is able to generalize according to multiple drug design exercises, including docking-specific measures of drug-likeness.

One notable VAE, generative tensorial reinforcement learning (GENTRL), was developed by Zhavoronkov et al. [80]. By compiling the techniques of RL, variational inference, and tensor decompositions, the VAE developed molecules with high synthesizability, bioactivity, and diversity values. While the compounds were not experimentally validated, quantum computational analyses confirmed the viability of the proposed conformations. Born et al. [81] described PaccMannRL, in which two VAEs are utilized, enabling the consideration of both molecular information and transcriptomic data. Even before encoding experimental data concerning anticarcinogens, the model exhibited a bias towards highly bioactive chemical agents. These findings support the use of transcriptomic data in de novo drug discovery with machine learning, which could be beneficial in developing cancer treatments specific to an individual's genetic profile. Arcidiacono and Koes [82] employed a VAE to encode 3D molecular graphs, introducing the possibility of VAE-informed 3D coordinate prediction. While Drotár et al. [83] also encoded molecules as 3D graphs, they chose a constrained graph VAE, departing from the traditional architectural standard. Their 3D graphs included fewer indications of compound radii, but the addition of data on ligand-protein complexes observed experimentally produced more viable drug-like molecules.

Schultz et al. [84] developed DarkChem, a VAE-based software, and generated novel antagonists of the NMDA receptor as proof of concept. Specifically, DarkChem is tailored to predict metabolomic properties, doing so more efficiently than first-principles simulations. Data obtained in silico and experimentally were amalgamated to train and refine the model, improving its predictive accuracy. VAE models in particular benefit from experimental results, as demonstrated earlier by Drotár et al. Chenthamarakshan et al. [85] launched CogMol, another VAE with the capacity to process target-specific information. After the AE was trained to design molecules with a high binding affinity for the SARS-CoV-2 protease, presented ligands were assessed through molecular docking simulations. These procedures revealed that most substances generated expressed high target selectivity, in part due to the model's capacity to consider multiple constraints. Target-specific fine--tuning is not required, making the model suitable for circumstances in which data availability is limited. Most recently, Li and Ghosh [86] fabricated SQ-VAE, a scalable quantum generative autoencoder, showing that the use of quantum computing can be advantageous in

small molecule design, although the model is less efficient. This drawback is discerned in most applications of quantum computing. If greater computational power and a broader time window are both available, SQ-VAE could be considered for practical use.

AAEs were also implemented in several experiments. Kadurin et al. [87] developed druGAN, the first AAE invented to generate molecular graphs. In this study, the researchers revisited an earlier AAE model that suffered from reconstruction error and poor discriminative power. The addition of a hyperparameter to measure sources of error in both the discriminator and the generator resolved these issues. Nevertheless, the absence of RL was acknowledged as an enduring source of inaccuracy. The authors proposed that integrating a multi-GAN pipeline could improve this deficiency. Polykovskiy et al. [61] observed that an entangled conditional AAE (ECAAE) could originate drug-like molecules of a higher quality than those generated by other AEs. The novel ECAAE, joint entanglement, predictive entanglement, combined entanglement, and disentanglement were juxtaposed, with the ECAAE model generating the greatest proportion of unique molecules. While the combined entanglement approach outperformed all alternative models in generating high logP values for continuous features, the novel ECAAE had the highest predictive accuracy for binding energy. Bai and Yin [60] noted that the pharmacological space of kinase inhibitors could be augmented by applying an AAE, deemed an Ensemble of PCM-AAE (EPA). Because the EPA algorithm exploits the chemical space in its entirety, it can be operated with fewer data points. Additionally, the model does not require data obtained from docking simulations, reducing the overall computational expense associated with training and executing the model. Liu and Bailey [88] introduced an innovative method of stacking models, enhancing the Bidirectional AAE architecture presented in earlier studies. One of the AEs is adapted to generate molecules according to basic features of drug-like compounds and gene expression information, while the other learns chemical properties and optimizes suggested molecules. Compared to the single, unidirectional AAE implemented by Kadurin et al., the bidirectional AAE identified a greater proportion of valid structures. Blaschke et al. elaborate on the various applications of AEs in de novo molecular design, contrasting their performances to evaluate their competency in predicting bioactivity [89]. An AAE trained to follow a Uniform distribution generated the greatest proportion of valid SMILES. Furthermore, an AAE trained with a Gaussian distribution performed almost as well as the Uniform AAE. Both AAE models promoted a larger proportion of valid molecules with chemical similarity to the training set. Although AAEs are currently less prevalent in the literature than VAEs, the conclusions of Blaschke et al. insinuate that they should not be neglected in model selection.

4.1.3. GAN

Of the three model architectures present in publications since 2017, GAN was the least common. Sanchez-Lengeling et al. invented ORGANIC, an extension of the established ORGAN architecture [90]. As the first GAN designed to generate molecular structures, ORGANIC was often utilized as a point of comparison in later studies. Non-valid molecules were frequently generated. This drawback was attributed to the roughness of the chemical space explored, mode collapse, and insufficient hyperparameter variation. The adversarial threshold neural computer (ATNC) presented by Putin et al. [91] was a drastic improvement to ORGANIC, generating 72%, as opposed to only 7%, valid SMILES strings. Internal Diversity Clustering (IDC), an original objective reward function, encouraged the generation of a diverse molecular library. The model's reduced susceptibility to mode collapse may be attributed to the IDC method. The same group would later release a reinforced adversarial neural computer (RANC) based on a GAN model and RL [92]. The RANC system incorporates a differentiable neural computer (DNC) neural network; because, like LSTM networks, DNCs benefit from an explicit memory bank, this technique can prevent mode collapse.

Prykhodko et al. [93] constructed LatentGAN, finding that the model could generate compounds different from those produced by an RNN. This suggests that the two models could be run in conjunction with one another to obtain a more diverse set of drug candidates. MolGAN, a prominent generative model for drug-like molecules and other organic substances alike, was created by De Cao and Kipf [94]. Almost 100% of generated compounds proved valid. Additionally, because MolGAN is both implicit and likelihood-free, the computational expense associated with likelihood-based models is avoided. Yet, because the structures with which the model was trained were limited to a maximum heavy atom count of nine, the generator is unlikely to suggest more complex drug molecules. Méndez-Lucio et al. [95] stacked conditional GANs and the Wasserstein GAN with gradient penalty to improve the quality of generated molecules. Both gene expression data and structural representations of small molecules were utilized to influence the generator. Compared to structures identified by traditional similarity search procedures, including Euclidean and cosine distance, original molecules more closely resembled known active compounds. The l-Wasserstein distance metric was calculated again by Pölsterl and Wachinger [96]. The model, termed ALMGIG, incorporated a likelihood-free network, building upon the strengths of MolGAN. However, the inference network incorporated into ALMGIG significantly reduced the rate of mode collapse. Multiple modes of the distribution were identified by the reformed model. Mol-CycleGAN, produced by Maziarka et al. [97] and inspired by MolGAN, was exploited to compose molecules with defined structural or physicochemical properties. Similarity and logP values were assessed as a constrained optimization task. According to a comparison between Mol-CycleGAN, a junction tree VAE, and a graph convolution policy network, the novel GAN model demonstrated the greatest rate of improvement, indicating that it is an ideal framework for molecular optimization problems. Recently, Jacobs and Maragoudakis have developed ASYNT-GAN, a model intended to effect molecules in accordance with their binding affinity to specific proteins [98]. The model was shown to generalize beyond the target data collected to train the model. While the suggested method of encoding 3D coordinates representative of the ligand complexed with a protein target, the model fails to generate point clouds without any noise. A final adaptation of MolGAN, a quantum GAN with a hybrid generator (QGAN-HG), was described by Li et al. [99]. It was found that, by reducing the depth of the neural network, issues like instability and vanishing gradients could be avoided. To compensate for the fact that more time is needed to execute quantum neural networks, the parameters used in the original MolGAN framework were reduced. This reduced form of the model showed favorable performance only when implemented as a QGAN.

4.1.4. Evolutionary algorithms

Yoshikawa et al. [100] implemented ChemGE, a grammatical evolution model. Hundreds of the generated molecules were validated in docking studies. Furthermore, the capacity to run several evolution simulations simultaneously represents an additional advantage of ChemGE. Another significant evolutionary model, EvoMol, was devised by Leguy et al. [101]. Various advances were made in the EvoMol study—notably, the capacity to construct molecular graphs independent of initial data. Although the model also composed many non-drug-like molecules in separate experiments, goal-directed synthesis was utilized to compose molecules with specified bioactivities. Nigam et al. [102] created JANUS, a parallel tempered genetic algorithm deep neural network. Parallel propagation enables the maintenance of two populations during each generation interval; while one probes the pharmacological space, the other selects for desired features. These populations are not isolated, meaning that member exchange between each selection of molecules is feasible. Compared to various VAEs, ChemTS, ORGAN, and MolecularRNN, JANUS obtained the greatest average best penalized logP values. Park et al. [103] formed FasterGTS, a genetic algorithm with an MCTS. The similarity of generated molecules to reference drugs was optimal for both FasterGTS and ChemTS. The authors state that JANUS appears to have been overfitted, which may imply that the conclusions drawn in the original study were not

explored thoroughly. PASITHEA [104], a model drafted based on inceptionism techniques and evolutionary models for de novo molecular design, differs from previous genetic algorithms in that it directly, rather than indirectly, explores the chemical space. As a consequence, it is possible to probe the model's interpretation of the chemical space explicitly. This property, reversibility, is not observed in VAE models, which suffer from deficient transparency.

4.1.5. Other models

Models were placed under the "Other" category if they were only raised in two or fewer publications. First, a Transformer neural network for de novo drug design based on amino acid sequence data was described by Grechishnikova [105]. Information about a target protein guided the generation of new molecular structures using a self-attention mechanism. Another Transformer algorithm established by Hu et al. applied a Transformer algorithm to the SARS-CoV-2 main protease [106]. Because standard self-attention considers all tokens in a sequence, it is not preferable in tasks involving the prediction of only the next token. The pipeline utilized generated molecules with attributes similar to those detected in experimental studies of approved drugs. Gao et al. developed a generative network complex (GNC), a model incorporating both an AE and an LSTM network [107]. Unlike the AAE architecture, no GAN pipeline is included in the GNC. Instead, the model is characterized by features of both a VAE and an RNN model. Similarity scores were evaluated with molecular fingerprints and pharmacophore analysis; produced molecules performed well on both metrics. Both similarity and density of the training data with respect to desired property values were emphasized as necessary components of the training and testing datasets. The final uncategorized model was L-net, developed by Li et al. [108] to generate molecules in three dimensions. In the initial study, L-net was applied with MCTS to discover inhibitors of ABL1 kinase. In the second [109], the same basic algorithm was run, with a few methodological adjustments. Instead of ABL1 kinase, the SARS-Cov-2 protease was targeted. The authors also introduced DeepLigBuilder, a model capable of directly generating 3D ligand structures.

Changes in the frequencies of the described model architectures are shown in Fig. 10. GANs, VAEs, AAEs, and LSTM networks were among the earliest attempted architectures proposed for the generation of drug-like molecules de novo. LSTM-RNN models were the most common models in 2017, 2020, and 2021. VAEs were utilized often in 2018. The majority of experimental papers published in 2019 included GANs. RNN-GRUs did not appear in any experimental articles released in 2021, and GANs presented less frequently as well. As established, GANs are generally more susceptible to mode collapse than other generative models, which may have influenced their reduced presence in the literature from 2021. The use of AAEs and evolutionary algorithms was inconsistent; the first AAEs were introduced in 2017 and 2018 but were not present in any studies from 2019 to 2020. They then emerged in multiple publications from 2021. Evolutionary algorithms came into use around 2018 but were not related in 2019. Granted, fewer experimental articles were observed overall in 2019, so these gaps might simply indicate a broader chronological trend. Because the Transformer and L-Net algorithms are relatively recent, no definitive conclusions can be made about their future using the limited data presently available. The LSTM-RNN approach is likely to remain conventional, while the relevance of RNN-GRUs and GANs may decline as other models emerge.

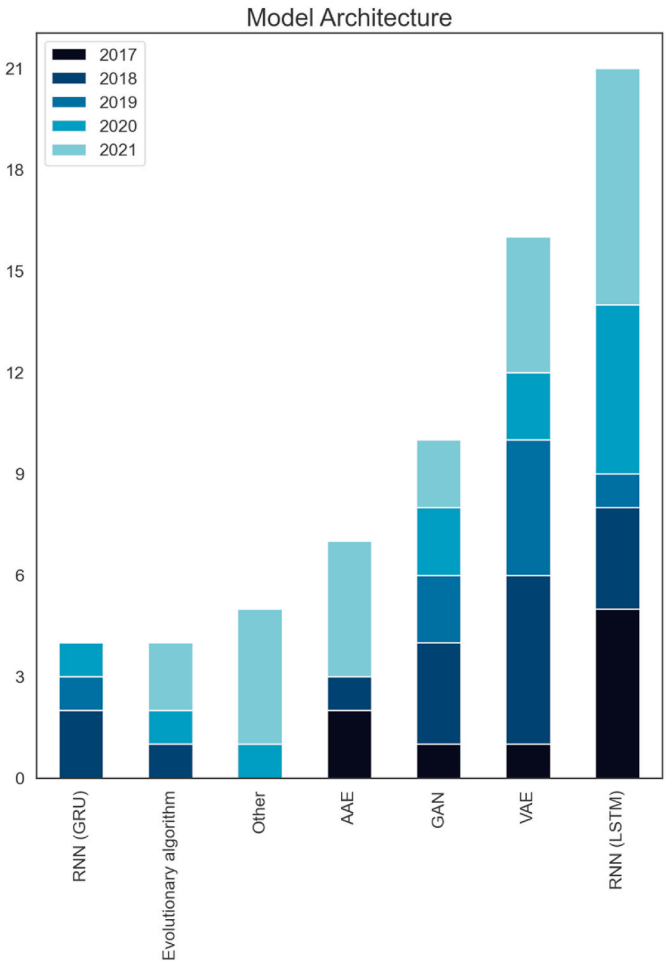

Fig. 10. The relative frequencies of the various machine learning model architectures observed in the review and their chronological progression.

4.2. Molecular representations

4.2.1. String-based representations

Of all molecular encoding options, canonical SMILES strings were the most frequent. Using SMILES frames the problem of de novo molecule generation as a natural language processing task. As such, the strings were usually either one-hot encoded or integer encoded, such that each atom or bond in a string could be formatted as a matrix. Most studies did not specify whether data was augmented, but eight of the 94 publications explicitly discussed methods to augment SMILES data. Training the encoder with enumerated SMILES, then decoding as canonical SMILES, results in the tightest clustering. Training the encoder or the decoder with enumerated SMILES yielded the best performance in QSAR modeling [56]. Additionally, generative models trained with enumerated SMILES produced a more diverse, unique, and valid collection of structures in less time than those that only included canonicalized SMILES [53,58,62,69,93,106,110,111].

SELFIES were also utilized to encode molecules, albeit less often than SMILES, possibly because SELFIES are comparatively new. SELFIES are advantageous in that they may serve as surjective representations of molecules. As a consequence, training a generative model with molecules encoded in the format eliminated the probability of generating invalid molecules [104]. SELFIES commonly delineated molecular structures in genetic algorithms [102,103], as the progressive modification of generated structures is more efficient with the use of a molecular representation that is entirely robust.

Of the models included in this review, only one, druGAN [87], was created to encode and generate molecular fingerprints. While other publications used molecular fingerprints to compare training data and the output, druGAN incorporated MACCS key fingerprints directly into the model. The authors conceded that "MACCS molecular fingerprints … are not [an] ideal representation of molecular structure."

4.2.2. 2D graphs

Molecules were represented with 2D graphs. To generate new SMILES strings, generative models must learn the rules of SMILES

syntax, some of which have no relevance to molecular structure [94,112]. The linearity of SMILES strings can also complicate the depiction of molecular properties, making goal-driven synthesis less straightforward [97]. Therefore, designating graphs to represent chemical structures in generative models could both reduce computational cost and equip algorithms to better process structural information. Relative to SMILES strings, 2D graphs are more suitable for RNNs [113]. Additionally, the use of 2D molecular graphs enables the sequential modification and validation of molecules [97,101]. 2D graphs can be encoded through a series of graph convolutions [76,97].

4.2.3. 3D graphs

Neither SMILES representations nor 2D graphs provide explicit details about a molecule's 3D conformation. Molecules may not be sufficiently characterized if only rendered as a function of their surface-level characteristics [114]. In scenarios where training data is limited, the machine learning algorithm will have few opportunities to learn implicit rules of molecular structure, such as bond order. By defining molecular coordinates and other physical rules of compound structure, the complications associated with low data can be mitigated [108]. 3D information might be encoded as a node-features and edge-weights matrix to specify the coordinates of atoms in 3D space [78,83]. Graph convolutional layers can integrate state encoding [108,109].

As shown in Fig. 11, the overwhelming majority of models are trained using non-randomized SMILES strings. Although new depictions, like SELFIES and 3D graphs, have become increasingly popular in recent years, SMILES are unlikely to disappear from the literature any time soon. They were the most common format across all the studied years, with the exception of 2019, during which more publications were found to have utilized 2D graphs. None of the experimental studies in 2021 explicitly mentioned SMILES randomization. Nevertheless, they may have been incorporated without any direct notice in the text. The appeal of canonical SMILES strings may be attributed to their prominent role in the history of cheminformatics. Recent approaches have yet to shape the field to such an extent, although they may be more effective in depicting molecules. 3D graphs seem to have attracted more interest in 2021, so their influence may continue beyond 2021. To optimize outcomes, the potential benefits of novel molecular encoding methods should not be neglected.

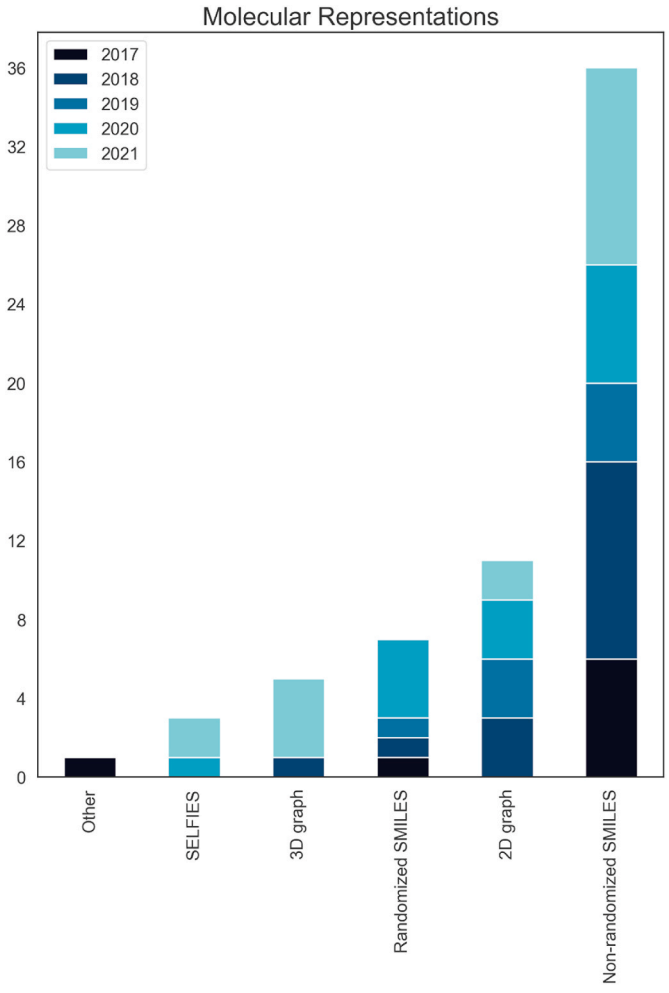

Fig. 11. The relative frequencies of the various molecular encoding methods observed in the review and their chronological progression.

4.3. Challenges

After reviewing the literature on the use of machine learning in de novo drug design, eight major challenges faced by developers were identified.

4.3.1. Diversity

In evaluating the viability of a model as a practical tool, molecular diversity was a common metric. To discover a drug candidate that demonstrates all selected properties, it is crucial to explore a wide range of options. Human chemists are aware of chemical diversity, but machine learning algorithms must be informed explicitly of diversity as a desired metric. Limited diversity may imply the issue of mode collapse, especially in GANs [87,90,94,115–118]. In other words, the model fails to generate a range of outputs, instead generating only one final molecule. Sometimes, a lack of diversity may indicate excessive similarity to the training data [119].

Solutions proposed to prevent uniformity in molecular structure involve either the alteration of the input data or the adjustment of the model architecture. An implicit, limited vocabulary of SMILES strings could prevent mode collapse [54]. Moreover, randomized SMILES strings have repeatedly proven beneficial in the generation of diverse outputs [53,56,62,64,110,111]. When a greater number of distinct SMILES tokens are provided as training data, the distribution from which a model may sample is wider, providing more opportunities for unique molecule generation [81]. For scaffold-based generative models, smaller scaffolds can introduce new pathways for structural diversification [79]. VAEs and GANs, which tend to suffer mode collapse more frequently than RNNs [71,115] may benefit from the representation of molecules not as SMILES strings but as molecular graphs [52,72,78,97], although this remedy is not foolproof [94].

Modifications can be made to the learning process itself to prevent mode collapse. Reinforcement and transfer learning have been tremendously successful techniques in guiding the generation of distinct molecules learning [65,66,71,75,78,119,120]. Although they are less powerful alone, AAE and GAN models can produce a wider variety of molecules when incorporated into ensemble models, as shown by Bai et al. [60]. Slight improvement in AAE output has also been observed when bidirectional functionality is enabled [88]. Because they integrate an explicit memory bank, DNCs are able to enhance the GAN operation [73,92]. In addition, GANs may benefit from an internal diversity clustering reward [91]. Nevertheless, because VAE and GAN models so frequently fail to generate a collection of diverse molecules, they should be used with caution. RNNs are less likely to produce uniform libraries of molecules. Even so, their performance can be further enhanced with special exploration strategies [74] or analog series design [58].

4.3.2. Limited data

Machine learning is dependent on adequate data. In this regard, new drug targets are particularly troublesome, as little, if any, training data is available upon their discovery [54,121,122]. Even if a large volume of assays were conducted on a single target, analyzing those assays requires

considerable resources, some of which may be inaccessible [17,123,124]. Furthermore, data quantity is not the only determinant of a model's success; an algorithm is only as good as the data with which it is trained. Although 3D molecular graphs can be generated in silico, deriving 3D molecular graphs directly from crystallographic data and encoding them as such might produce more accurate synthetic structures [108]. When optimizing a specific characteristic, training the model on a dataset with molecules exhibiting a wide range of property values is necessary [107]. Using one model in isolation is ineffective under these circumstances.

Aware of the limits of assay data, many have refined models able to use details about the amino acid or nucleotide sequence of a target to propose relevant hit-like molecules [15,66,95,105]. Models capable of transfer learning or active learning are robust, despite the scarcity of data, as they reduce the amount of manual annotation required of human users and automatically identify meaningful knowledge gaps [16,58,120]. SMILES enumeration, beyond improving output diversity, can serve as a form of data augmentation [56,62,69,106]. There is evidence that unsupervised pre-training enhances model performance on small datasets [87]. Some models, like EvoMol, were contrived to function independent of starting data [101]. GENTRL was also shown to operate with incomplete information [80]. Even if none of these solutions are appropriate for a particular model, the use of generative machine learning in conjunction with other algorithms and domain knowledge would still expedite the drug discovery process and enable experts to explore the chemical space in silico [59,93].

4.3.3. Interpretability

Designing a model that is both interpretable and effective has proven challenging [61,100,125,126]. This problem relates to "black box" machine learning algorithms, which are only understood in terms of their inputs and outputs. VAE and GAN models often suffer from this weakness, as they rely on non-linear operations [123], though other models may similarly exhibit this characteristic [105]. In the interest of interpretable machine learning, integrated gradients, graph convolutions, sensitivity analysis, and variable importance have been proposed [70,124]. Evolutionary algorithms allow users to visualize the exploration of chemical space [101] rather than reducing the contribution of the model to its immediate product. The process of molecular graph generation provides valuable insight into attempted structural components. Another technique entails the inversion of the internal molecular representation, as described by Shen et al., the developers of PASITHEA [104]. They found that the model learned to associate nitrogen atoms with lower logP values, using this knowledge to optimize generated molecules accordingly. Similar techniques might be opportune in determining the structural basis for chemical properties.

4.3.4. Synthesizability

Another issue that complicates the practical integration of machine learning in the laboratory is the tendency of algorithms to generate molecules that are difficult to synthesize [127,128]. In particular, VAEs regularly fail to generate synthesizable molecules [74,84,115]. Goal-directed generation tasks are especially troublesome, producing a greater proportion of insynthesizable molecules even when biased by training set synthesizability [129]. However, multiple algorithms reviewed in this study were organized with this ubiquitous problem in mind [68], designating synthesizability as a metric to evaluate model quality. Machine learning algorithms informed of retrosynthesis pathways were more competent in assigning weights to structural components [80,85,122,130–132]. Ghiandoni et al. introduced a multi-label reaction class recommendation algorithm to restrict molecule generation to synthesizable compounds [133]. Including metrics of synthetic accessibility either within the reward function or externally could also reduce the proportion of synthetically unfeasible molecules generated [51,52,55,68]. Another benefit of SMILES randomization could manifest in higher synthesizability scores [110].

4.3.5. Comparability

Many studies raised the question of model generalizability and comparability, arguing that both benchmark datasets and universal measures of model refinement should be established [126]. Percent validity of generated molecules, uniqueness, and diversity were common metrics used to assess model performance, but these standards alone may be an insufficient means of comparing models [96,117]. Different encodings further complicate the issue, as some models adapt to string-based representations well, while others work better for molecular graphs [60,113]. Databases frequently demonstrate distributions that are non-identical and dependent [134]. While the ChEMBL and ZINC databases contain a wide range of drug-like molecules, their versatility prevents them from acting as standardized inputs.

Since their publication, benchmark datasets MOSES [135] and GuacaMol [136] have become widespread, especially with respect to de novo drug design tasks [111,137]. In 2018, Preuer et al. proposed Fréchet ChemNet distance (FCD) as a metric for the quality of models produced by machine learning algorithms [138]. FCD is more robust than ambiguous measures of model performance such as validity, uniqueness, novelty, and KL divergence [134]. Regardless, because drug design is a complex process that must be executed well in three dimensions, novel compounds should also be validated with docking-based benchmarks [139]. As these metrics and benchmarks come into general use, the obstacles associated with comparing the performance of generative models will be mitigated.

4.3.6. Multi-objective capability

Here, multi-objective capability definitively refers to the interaction between generated molecules and ADMET properties. Avoiding the synthesis of toxic or otherwise undesirable compounds depends on the capacity to optimize a molecule in several contexts [15,16,81,140–142]. Moreover, certain tasks in drug design benefit from molecules that are bioactive relative to multiple targets. Models unable to account for the numerous factors imperative to the discovery of compounds with desired activities are less feasible. Even if the originated molecules are assessed for their properties via external models, this technique can come at a high computational cost.

As such, model-intrinsic selection of desirable attributes could empower the selection of structures with multiple selected activities without requiring extrinsic validation [68,85]. Similarly, conditional models are more efficient in generating compounds with appropriate features [75,112]. Both solutions involve the sequential selection of specific structural elements, as do evolutionary algorithms [103]. Evolutionary operators can monitor compound fitness at various stages in the generative process, so compounds found to be inactive or toxic can be removed from the population before they are suggested. When enough data is available to do so, fine-tuning models to favor multi-target compounds may result in a greater volume of multi-target drug candidates. Ensuring accurate multi-objective predictions also depends on the consistency of such values across several models, which can be enhanced by ensemble AE-LSTM architectures [107]. Other proposed countermeasures include nondominated sorting [143] and improved uncertainty quantification [134].

4.3.7. Molecule size

While smaller molecules tend to demonstrate reduced synthetic complexity, certain drug targets require larger molecules. This restricts the use of many generative algorithms, as several are capable only of generating very small molecules [77,82,83]. QM9, an earlier dataset used to train some de novo drug design algorithms, contains molecules with up to 9 heavy atoms, limiting the diversity and efficacy of generated compounds as drug-like substances [74,94]. Even with access to datasets like MOSES and GuacaMol, which contain larger, more drug-like molecules, the computational expense of encoding and decoding molecular structures represented in 3D or as quantum systems can be too great when handling large molecules [86].

It is important to note that the complex nature of drug design does not entail that cheminformatics algorithms must be complex by definition. In some cases, simpler models can operate both more efficiently and more effectively [62]. Adjustments to the organization of memory in an algorithm have proven effective as well; for example, by using a stacked memory or a DNC, longer compounds may be included more frequently, thus diversifying the effected library [51,92]. Logically, including a wider variety of input SMILES characters in the training set gives the model more freedom in exploiting the chemical space [81].

4.3.8. Uncertainty

The problem of uncertainty estimation and active learning implementation was rarely mentioned, but the importance of accounting for uncertainty is necessary for designing compounds with conviction. As previously stated, accuracy alone cannot justify the use of machine learning to inform human decision-making. If machine learning is to become a reliable guide of drug discovery, it will be essential to incorporate more precise means of estimating uncertainty [124,134]. Because uncertainty has received less attention than the other challenges disclosed in this review, fewer solutions have been raised. For generative tasks, Venn-ABERS predictors and Bayesian techniques of uncertainty estimation are considered some of the most appropriate [123,134].

In Fig. 12, the frequency of these issues in the literature for each year is presented. In 2021, the number of publications that indicated the problem of limited diversity decreased, suggesting that the models developed in earlier initiatives were successful in managing the issue.

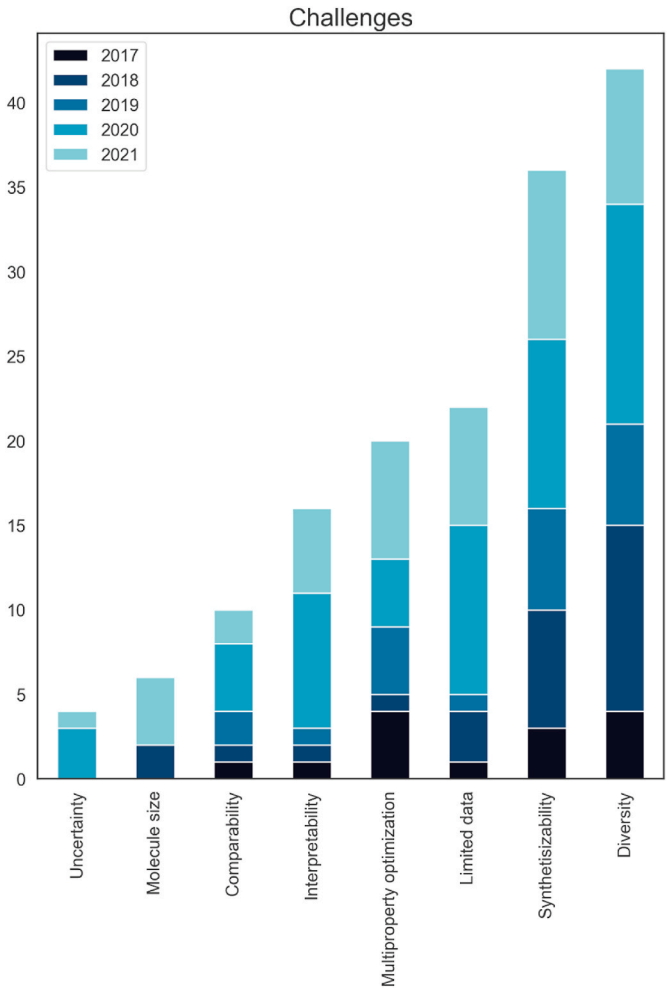

Fig. 12. The relative frequencies of the various challenges in de novo drug design using machine learning observed in the review and their chronological progression.

On the contrary, synthesizability was addressed proportionately over the past five years. It has remained a consistent issue. In 2021, synthesizability and multi-property optimization were the most frequently referenced obstacles to de novo drug design employing machine learning. Diversity, synthesizability, data quantity, and interpretability were relatively common subjects in the 2020 literature, occurring at similar frequencies. In the three years prior, diversity and synthesizability were the most prevalent subjects of interest. While those challenges appearing in the literature since 2017 were some of the first, interpretability, molecule size, and uncertainty emerged more recently. Discussion of the size of generated molecules was inconsistent, beginning in 2018 only to emerge again in 2021. Each topic continued to recur in 2021, so it is possible that studies published in 2022 will demonstrate similar trends.

Addressing these difficulties is evidently dependent not only on broad methodological reforms but also on model selection. To summarize, GAN frameworks are prone to limited diversity, considering their vulnerability to mode collapse. Mitigating these challenges may involve the use of IDC, an inference network, or of a DNC network. Furthermore, GANs can be more efficient than other generative models, as it is feasible to construct them without likelihood-based functions. Similarly, RNN networks have proven successful in general, exhibiting efficiency superior to that of VAEs. They have been applied to mitigate deficient synthetic feasibility, diversity, and multi-property optimization. RL and TL are crucial factors in augmenting diversity and synthetic feasibility. Comparative studies of GRU- and LSTM-RNNs suggest that LSTM networks are generally preferable concerning metrics of uniqueness, validity, and similarity. AEs, including VAEs and AAEs, represent a considerable proportion of generative algorithms involved in de novo drug design. RNNs and VAEs can both infer chemical laws without the explicit definition of these guidelines or of target-specific data. An obstacle frequently encountered in VAE models is the restricted size of generated compounds. Therefore, VAEs should be employed principally when simple compounds are desired. They are useful in instances of limited data, as several have functioned without target-precise data or with incomplete information about structures and properties. As for AAEs, entanglement can also improve model operation in circumstances of low data. Uniform, as opposed to Gaussian, distribution is ideal for AAE training. AAE networks are not the only available ensemble processes; the recent introduction of the GNC architecture provides meaningful insight regarding the value of integrating multiple models to guide de novo drug design, rather than relying on a single framework to model all relevant properties accurately. Evolutionary algorithms like EvoMol and PASITHEA are highly interpretable. This quality is necessary for facilitating the interdisciplinary advancement of de novo drug design methods. Finally, the new class of Transformer models has recently shown promise in expanding interpretability and synthesizability.

5. Conclusions

De novo drug design is one of many domains in which machine learning could inform human decision-making. Because machine learning algorithms are able to automate the recognition and prioritization of statistical patterns, they could also provide valuable new details about molecular structure and activity. In this systematic review, original research papers and review articles from the past five years were applied to extract information about trends in generative models for de novo drug design, challenges associated with the task, and formats utilized to encode data on chemical structure. LSTM-RNN, GAN, and VAE architectures appeared most frequently in the literature overall, but recent algorithms have yielded equally promising results. Eight major areas of concern were identified: the diversity, or lack thereof, of generated molecular libraries; the synthesizability of generated molecules; the limited quantity of assay data; model interpretability; the demand for multi-property optimization; the ability to compare model

performances; molecule size constraints; and measures of uncertainty in model evaluation. LSTM-RNN networks are optimized when combined with RL or TL. GANs require additional architecture support to prevent mode collapse. Adaptations of VAEs enabling the generation of more complex molecules and low-data learning were particularly successful. Mitigating errors related to these global problems requires adjustments to both the frameworks themselves and input adjustment. String-based representations of molecules were the most common, with SMILES being the most widespread format. 2D and 3D molecular graphs have also shown promise. Current methods may be expanded upon, and there is plenty of room for improvement. In the years to come, the exciting, rapid advancement of machine learning-aided drug discovery should be followed closely as an application of artificial intelligence in medicine.

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

No funding was received in conducting this systematic review.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.compbiomed.2022.105403.

References

[1] W. Duch, et al., Artificial intelligence approaches for rational drug design and discovery, Curr. Pharmaceut. Des. 13 (14) (2007) 1497–1508, https://doi.org/10.2174/138161207780765954.
[2] A. Lima, et al., Use of machine learning approaches for novel drug discovery, Expet Opin. Drug Discov. 11 (3) (2016) 225–239, https://doi.org/10.1517/17460441.2016.1146250.
[3] F. Zhu, et al., Clinical success of drug targets prospectively predicted by in silico study, Trends Pharmacol. Sci. 39 (3) (2018) 229–231, https://doi.org/10.1016/j.tips.2017.12.002.
[4] Y. Li, et al., Clinical trials, progression-speed differentiating features and swiftness rule of the innovative targets of first-in-class drugs, Brief, Bioinformatics 21 (2) (2020) 649–662, https://doi.org/10.1093/bib/bby130.
[5] R. Gómez-Bombarelli, et al., Design of efficient molecular organic light-emitting diodes by a high-throughput virtual screening and experimental approach, Nat. Mater. 15 (2016), https://doi.org/10.1038/nmat4717, 1120–11.
[6] T. Scior, et al., Recognizing pitfalls in virtual screening: a critical review, J. Chem. Inf. Model. 52 (4) (2012) 867–881, https://doi.org/10.1021/ci200528d.
[7] D. Rumelhart, et al., Learning representations by back-propagating errors, Nat 323 (6088) (1986) 533–536, https://doi.org/10.1038/323533a0.
[8] I. Goodfellow, et al., Generative adversarial nets, in: Adv. Neural Inf. Process. Syst., 27, 2014, in: https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf.
[9] D.P. Kingma, M. Welling, Auto-encoding Variational Bayes, 2013 *arXiv preprint arXiv:1312.6114*.
[10] K. Wang, X. Wan, Automatic generation of sentimental texts via mixture adversarial networks, Artif. Intell. 275 (2019) 540–558, https://doi.org/10.1016/j.artint.2019.07.003.
[11] O. Lopez-Rincon, et al., Algorithmic music composition based on artificial intelligence: a survey, IEEE Int. Conf. Electron. Comput. Commun. Technol. (2018) 187–193, https://doi.org/10.1109/CONIELECOMP.2018.8327197.
[12] J. Yu, et al., Generative Image Inpainting with Contextual Attention, CPVR, 2018, pp. 5505–5514, https://doi.org/10.48550/arXiv.1801.07892.
[13] B. Ivanovic, et al., Multimodal deep generative models for trajectory prediction: a conditional variational autoencoder approach, IEEE Robot. Autom. 6 (2) (2021) 295–302, https://doi.org/10.1109/LRA.2020.3043163.
[14] T. Unterthiner, et al., Deep learning as an opportunity in virtual screening, NeurIPS 27 (2014) 1–9.
[15] A. Bender, I. Cortés-Ciriano, Artificial intelligence in drug discovery: what is realistic, what are illusions? Part 1: ways to make an impact, and why we are not there yet, Drug Discov. Today 26 (2) (2021) 511–524, https://doi.org/10.1016/j.drudis.2020.12.009.
[16] M. Elbadawi, et al., Advanced machine-learning techniques in drug discovery, Drug Discov. Today 26 (3) (2021) 769–777, https://doi.org/10.1007/s11030-021-10217-3.
[17] D. Paul, et al., Artificial intelligence in drug discovery and development, Drug Discov. Today 26 (1) (2021) 80, https://doi.org/10.1016/j.drudis.2020.10.010.
[18] A. Makhzani, et al., Adversarial Autoencoders, 2015 arXiv preprint arXiv:1511.05644.
[19] P.A. Vikhar, Evolutionary Algorithms: A Critical Review and its Future Prospects, ICGTSPICC., 2016, pp. 261–265, https://doi.org/10.1109/ICGTSPICC.2016.7955308.
[20] D. Weininger, SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules, J. Chem. Inf. Comput. Sci. 28 (1) (1988) 31–36, https://doi.org/10.1021/ci00057a005.
[21] M. Krenn, et al., Self-Referencing Embedded Strings (SELFIES): a 100% robust molecular string representation, Mach. Sci. Technol. 1 (4) (2020), 045024, https://doi.org/10.1088/2632-2153/aba947.
[22] A. Cereto-Massagué, et al., Molecular fingerprint similarity search in virtual screening, Methods 71 (2015) 58–63, https://doi.org/10.1016/j.ymeth.2014.08.005.
[23] D. Rogers, M. Hahn, Extended-connectivity fingerprints, J. Chem. Inf. Model. 50 (5) (2010) 742–754, https://doi.org/10.1021/ci100050t.
[24] S.E. Ward, P. Beswick, What does the aromatic ring number mean for drug design? Expet Opin. Drug Discov. 9 (9) (2014) 995–1003, https://doi.org/10.1517/17460441.2014.932346.
[25] L. David, et al., Molecular representations in AI-driven drug discovery: a review and practical guide, J. Cheminf. 12 (2020) 56, https://doi.org/10.1186/s13321-020-00460-5.
[26] E. Estrada, Characterization of 3D molecular structure, Chem. Phys. Lett. 319 (5–6) (2000) 713–718, https://doi.org/10.1016/S0009-2614(00)00158-5.
[27] H. Cho, I.S. Choi, Three-dimensionally Embedded Graph Convolutional Network (3dgcn) for Molecule Interpretation, 2018, https://doi.org/10.1002/cmdc.201900458 *arXiv preprint arXiv:1811.09794*.
[28] S. Liu, et al., Pre-training Molecular Graph Representation with 3D Geometry, 2021 arXiv preprint arXiv:2110.07728.
[29] J.J. Irwin, B.K. Shoichet, ZINC–a free database of commercially available compounds for virtual screening, J. Chem. Inf. Model. 45 (1) (2005) 177–182, https://doi.org/10.1021/ci049714+.
[30] D.S. Wishart, et al., DrugBank 5.0: a major update to the DrugBank database for 2018, Nucleic Acids Res 46 (D1) (2018) D1074–D1082, https://doi.org/10.1093/nar/gkx1037.
[31] J.F. Armstrong, et al., The IUPHAR/BPS Guide to PHARMACOLOGY in 2020: extending immunopharmacology content and introducing the IUPHAR/MMV Guide to MALARIA PHARMACOLOGY, Nucleic Acids Res 48 (D1) (2020), https://doi.org/10.1093/nar/gkz951. D1006–D1021.
[32] A. Gaulton, et al., The ChEMBL database in 2017, Nucleic Acids Res 45 (D1) (2017) D945–D954, https://doi.org/10.1093/nar/gkw1074.
[33] M.H. Saier, et al., The transporter classification database (TCDB): 2021 update, Nucleic Acids Res 49 (D1) (2021), https://doi.org/10.1093/nar/gkaa1004. D461–D467.
[34] T. Fu, et al., VARIDT 2.0: structural variability of drug transporter, Nucleic Acids Res 50 (D1) (2022), https://doi.org/10.1093/nar/gkab1013. D1417–D1431.
[35] G. Klebe, et al., Molecular similarity indices in a comparative analysis (CoMSIA) of drug molecules to correlate and predict their biological activity, J. Med. Chem. 37 (24) (1994) 4130–4146, https://doi.org/10.1021/jm00050a010.
[36] A. Bender, R.C. Glen, Molecular similarity: a key technique in molecular informatics, Org. Biomol. Chem. 2 (2004) 3204–3218, https://doi.org/10.1039/B409813.
[37] J. Bueno, ADMETox: Bringing Nanotechnology Closer to Lipinski's Rule of Five, Preclinical Evaluation of Antimicrobial Nanodrugs, Springer, Cham, 2020, pp. 61–74.
[38] A. Daina, et al., SwissADME: a free web tool to evaluate pharmacokinetics, drug-likeness and medicinal chemistry friendliness of small molecules, Sci. Rep-UK 7 (1) (2017) 1–13, https://doi.org/10.1038/srep42717.
[39] S. Saikia, M. Bordoloi, Molecular docking: challenges, advances and its use in drug discovery perspective, Curr. Drug Targets 20 (5) (2019) 501–521, https://doi.org/10.2174/1389450119666181022153016.
[40] A.E. Lohning, et al., A practical guide to molecular docking and homology modelling for medicinal chemists, Curr. Top. Med. Chem. 17 (18) (2017) 2023–2040, https://doi.org/10.2174/1568026617666170130110827.
[41] W. Xue, et al., What contributes to serotonin-norepinephrine reuptake inhibitors' dual-targeting mechanism? The key role of transmembrane domain 6 in human serotonin and norepinephrine transporters revealed by molecular dynamics simulation, ACS Chem. Neurosci. 9 (5) (2018) 1128–1140, https://doi.org/10.1021/acschemneuro.7b00490.
[42] W. Xue, et al., Computational identification of the binding mechanism of a triple reuptake inhibitor amitifadine for the treatment of major depressive disorder, Phys. Chem. Chem. Phys. 20 (9) (2018) 6606–6616, https://doi.org/10.1039/c7cp07869b.
[43] T. Fu, et al., Subtype-selective mechanisms of negative allosteric modulators binding to group I metabotropic glutamate receptors, Acta Pharmacol. Sin. 42 (8) (2021) 1354–1367, https://doi.org/10.1038/s41401-020-00541-z.
[44] D. Moher, et al., Preferred reporting items for systematic review and meta-analysis protocols (PRISMA-P) 2015 statement, Syst. Rev. 4 (1) (2015) 1, https://doi.org/10.1186/2046-4053-4-1.
[45] M. Hartenfeller, et al., DOGS: reaction-driven de novo design of bioactive compounds, PLoS Comput. Biol. 8 (2) (2012), e1002380, https://doi.org/10.1371/journal.pcbi.1002380.
[46] R. Gómez-Bombarelli, et al., Automatic chemical design using a data-driven continuous representation of molecules, ACS Cent. Sci. 4 (2) (2018) 268–276, https://doi.org/10.1021/acscentsci.7b00572.

[47] M. Segler, et al., Generating focused molecule libraries for drug discovery with recurrent neural networks, ACS, Cent. Sci. 4 (1) (2018) 120–131, https://doi.org/10.1021/acscentsci.7b00512.

[48] Critical Appraisal Skills Programme, CASP Qualitative Studies Checklist, 2019. https://casp-uk.net/. (Accessed 20 December 2021). Accessed.

[49] Critical Appraisal Skills Programme, CASP Systematic Review Checklist, 2019. https://casp-uk.net/. (Accessed 20 December 2021). Accessed.

[50] M. Olivecrona, et al., Molecular de-novo design through deep reinforcement learning, J. Cheminf. 9 (2017) 48, https://doi.org/10.1186/s13321-017-0235-x.

[51] M. Popova, et al., Deep reinforcement learning for de novo drug design, Sci. Adv. 4 (7) (2018) 7885, https://doi.org/10.1126/sciadv.aap7885.

[52] X. Yang, et al., ChemTS: an efficient python library for *de novo* molecular generation, Sci. Technol. Adv. Mater. 18 (1) (2014) 972–976, https://doi.org/10.1080/14686996.2017.1401424.

[53] P. Ertl, et al., Silico Generation of Novel, Drug-like Chemical Matter Using the LSTM Neural Network, 2017 arXiv preprint arXiv:1712.07449.

[54] A. Gupta, et al., Generative recurrent networks for de novo drug design, Mol. Inform. 37 (1–2) (2018) 1700111, https://doi.org/10.1002/minf.201700111.

[55] D. Merk, et al., De novo design of bioactive small molecules by artificial intelligence, Mol. Inform. 37 (1–2) (2018) 1700153, https://doi.org/10.1002/minf.201700153.

[56] E.J. Bjerrum, B. Sattarov, Improving chemical autoencoder latent space and molecular *de novo* generation diversity with heteroencoders, Biomolecules 8 (4) (2018) 131, https://doi.org/10.3390/biom8040131.

[57] D. Bruns, et al., Synthetic activators of cell migration designed by constructive machine learning, ChemistryOpen 8 (10) (2019) 1303–1308, https://doi.org/10.1002/open.201900222.

[58] D. Yonchev, J. Bajorath, DeepCOMO: from structure-activity relationship diagnostics to generative molecular design using the compound optimization monitor methodology, J. Comput. Aided Mol. Des. 34 (2020) 1207–1218, https://doi.org/10.1007/s10822-020-00349-3.

[59] P. Schneider, et al., Virtual screening and design with machine intelligence applied to pim-1 kinase inhibitors, Mol. Inform. 39 (9) (2020), e2000109, https://doi.org/10.1002/minf.202000109.

[60] X. Bai, Y. Yin, Exploration and augmentation of pharmacological space via adversarial auto-encoder model for facilitating kinase-centric drug development, J. Cheminf. 13 (2021) 95, https://doi.org/10.1186/s13321-021-00574-4.

[61] D. Polykovskiy, et al., Entangled conditional adversarial autoencoder for de Novo drug discovery, Mol. Pharm. 15 (10) (2018) 4398–4405, https://doi.org/10.1021/acs.molpharmaceut.8b00839.

[62] X. Li, et al., Chemical space exploration based on recurrent neural networks: applications in discovering kinase inhibitors, J. Cheminf. 12 (2020) 42, https://doi.org/10.1186/s13321-020-00446-3.

[63] P. Schneider, et al., Virtual screening and design with machine intelligence applied to pim-1 kinase inhibitors, Mol. Inform. 39 (9) (2020) 2000109, https://doi.org/10.1002/minf.202000109.

[64] M.V.S. Santana, F.P. Silva, De novo design and bioactivity prediction of SARS-CoV-2 main protease inhibitors using recurrent neural network-based transfer learning, BMC Chem 15 (2021) 8, https://doi.org/10.1186/s13065-021-00737-2.

[65] S. Amilpur, R. Bhukya, Predicting novel drug candidates against Covid-19 using generative deep neural networks, J. Mol. Graph. Model. 110 (2022), https://doi.org/10.1016/j.jmgm.2021.108045, 108045.

[66] M. Thomas, et al., Comparison of structure- and ligand-based scoring functions for deep generative models: a GPCR case study, J. Cheminf. 13 (2021) 39, https://doi.org/10.1186/s13321-021-00516-0.

[67] T. Blaschke, J. Bajorath, Fine-tuning of a generative neural network for designing multi-target compounds, J. Comput. Aided Mol. Des. (2021) 1–9, https://doi.org/10.1007/s10822-021-00392-8.

[68] M. Moret, et al., Beam search for automated design and scoring of novel ROR ligands with machine intelligence, Angew. Chem. Int. Ed. 60 (35) (2021) 19477–19482, https://doi.org/10.1002/anie.202104405.

[69] R. van Deursen, et al., GEN: highly efficient SMILES explorer using autodidactic generative examination networks, J. Cheminf. 12 (2020) 22, https://doi.org/10.1186/s13321-020-00425-8.

[70] Q. Bai, et al., Application advances of deep learning methods for de novo drug design and molecular dynamics simulation, Wires. Comput. Mol. Sci. (2021) e1581, https://doi.org/10.1002/wcms.1581.

[71] S. Zheng, et al., QBMG: quasi-biogenic molecule generator with deep recurrent neural network, J. Cheminf. 11 (2019) 5, https://doi.org/10.1186/s13321-019-0328-9.

[72] M. Popova, MolecularRNN: Generating Realistic Molecular Graphs with Optimized Properties, 2019 *arXiv preprint arXiv:1905.13372*.

[73] T. Blaschke, et al., Memory-assisted reinforcement learning for diverse molecular de novo design, J. Cheminf. 12 (2020) 68, https://doi.org/10.1186/s13321-020-00473-0.

[74] X. Liu, et al., An exploration strategy improves the diversity of de novo ligands using deep reinforcement learning: a case for the adenosine A2A receptor, J. Cheminf. 11 (2019) 35, https://doi.org/10.1186/s13321-019-0355-6.

[75] J. Lim, et al., Molecular generative model based on conditional variational autoencoder for de novo molecular design, J. Cheminf. 10 (2018) 31, https://doi.org/10.1186/s13321-018-0286-7.

[76] R. Assouel, et al., Defactor: Differentiable Edge Factorization-Based Probabilistic Graph Generation, 2018 *arXiv preprint arXiv:1811.09766*.

[77] M. Simonovsky, N. Komodakis Graphvae, Towards generation of small graphs using variational autoencoders, Int. Conf. Artif. Neural Netw. (2018) 412–422, https://doi.org/10.48550/arXiv.1802.03480. Springer, Cham.

[78] B. Samanta, et al., Nevae: a deep generative model for molecular graphs, J. Mach. Learn. Res. 21 (114) (2021) 1–33, https://doi.org/10.48550/arXiv.1802.05283.

[79] Y. Li, et al., Deepscaffold: a comprehensive tool for scaffold-based de novo drug discovery using deep learning, J. Chem. Inf. Model. 60 (1) (2019) 77–91, https://doi.org/10.1021/acs.jcim.9b00727.

[80] A. Zhavoronkov, et al., Deep learning enables rapid identification of potent DDR1 kinase inhibitors, Nat. Biotechnol. 37 (9) (2019) 1038–1040, https://doi.org/10.1038/s41587-019-0224-x, 2019.

[81] J. Born, et al., PaccMannRL: de novo generation of hit-like anticancer molecules from transcriptomic data via reinforcement learning, iScience 24 (4) (2021) 102269, https://doi.org/10.1016/j.isci.2021.102269.

[82] M. Arcidiacono, D.R. Koes, MOLUCINATE: A Generative Model for Molecules in 3D Space, 2021 *arXiv preprint arXiv:2109.15308*.

[83] P. Drotár, Structure-aware Generation of Drug-like Molecules, 2021 *arXiv preprint arXiv:2111.04107*.

[84] K.J. Schultz, Application and assessment of deep learning for the generation of potential NMDA receptor antagonists, Phys. Chem. Chem. Phys. 23 (2021) 1197–1214, https://doi.org/10.1039/D0CP03620J.

[85] V. Chenthamarakshan, Cogmol: Target-specific and Selective Drug Design for COVID-19 Using Deep Generative Models, 2020 *arXiv preprint arXiv*:2004.01215.

[86] J. Li, S. Ghosh, Scalable Variational Quantum Circuits for Autoencoder-Based Drug Discovery, 2021 *arXiv preprint arXiv:2112.12563*.

[87] A. Kadurin, et al., DruGAN: an advanced generative adversarial autoencoder model for de Novo generation of new molecules with desired molecular properties in silico, Mol. Pharm. (2017) 14, https://doi.org/10.1021/acs.molpharmaceut.7b00346.

[88] Y. Liu, J. Bailey, De Novo Molecular Generation with Stacked Adversarial Model, 2021 *arXiv preprint arXiv:2110.12454*.

[89] T. Blaschke, et al., Application of generative autoencoder in de novo molecular design, Mol. Inform. 37 (1–2) (2018) 1700123, https://doi.org/10.1002/minf.201700123.

[90] B. Sanchez-Lengeling, et al., Optimizing Distributions over Molecular Space. An Objective-Reinforced Generative Adversarial Network for Inverse-Design Chemistry (ORGANIC), 2017, https://doi.org/10.26434/chemrxiv.5309668.v3. *ChemRxiv*.

[91] E. Putin, et al., Adversarial threshold neural computer for molecular de novo design, Mol. Pharm. 15 (2018), https://doi.org/10.1021/acs.molpharmaceut.7b01137.

[92] E. Putin, et al., Reinforced adversarial neural computer for de novo molecular design, J. Chem. Inf. Model. 58 (6) (2018) 1194–1204, https://doi.org/10.1021/acs.jcim.7b00690.

[93] O. Prykhodko, et al., A de novo molecular generation method using latent vector based generative adversarial network, J. Cheminf. 11 (2019) 74, https://doi.org/10.1186/s13321-019-0397-9.

[94] N. De Cao, T. Kipf, MolGAN: an Implicit Generative Model for Small Molecular Graphs, 2018 *arXiv preprint arXiv:1805.11973*.

[95] O. Méndez-Lucio, et al., De novo generation of hit-like molecules from gene expression sigNat.s using artificial intelligence, Nat. Commun. 11 (2020) 10, https://doi.org/10.1038/s41467-019-13807-w.

[96] S. Pölsterl, C. Wachinger, Adversarial Learned Molecular Graph Inference and Generation, ECML PKDD 2020 p. 173-189. Springer, Cham. https://doi.org/10.1007/978-3-030-67661-2_11.

[97] Ł. Maziarka, et al., Mol-CycleGAN: a generative model for molecular optimization, J. Cheminf. 12 (2020) 2, https://doi.org/10.1186/s13321-019-0404-1.

[98] I. Jacobs, M. Maragoudakis, De novo drug design using artificial intelligence applied on SARS-CoV-2 viral proteins ASYNT-GAN, BioChem 1 (1) (2021) 36–48, https://doi.org/10.3390/biochem1010004.

[99] J. Li, et al., Quantum generative models for small molecule drug discovery, IEEE Trans. Quantum Eng. 2 (2021) 1–8 3103308, https://doi.org/10.1109/TQE.2021.3104804.

[100] N. Yoshikawa, et al., Population-based de novo molecule generation, using grammatical evolution, Chem. Lett. 47 (11) (2018) 1431–1434, https://doi.org/10.1246/cl.180665.

[101] J. Leguy, et al., EvoMol: a flexible and interpretable evolutionary algorithm for unbiased de novo molecular generation, J. Cheminf. 12 (2020) 55, https://doi.org/10.1186/s13321-020-00458-z.

[102] A. Nigam, et al., JANUS: Parallel Tempered Genetic Algorithm Guided by Deep Neural Networks for Inverse Molecular Design, 2021 arXiv preprint arXiv:2106.04011.

[103] S. Park, H. Lee, A Molecular Generative Model with Genetic Algorithm and Tree Search for Cancer Samples, 2021 *arXiv preprint arXiv*:2112.08959.

[104] C. Shen, et al., Deep Molecular Dreaming: inverse machine learning for de-novo molecular design and interpretability with surjective representations, Mach. Learn.: Sci. Technol. 2 (2021) 3, https://doi.org/10.1088/2632-2153/ac09d6.

[105] D. Grechishnikova, Transformer neural network for protein-specific de novo drug generation as a machine translation problem, Sci. Rep. 11 (2021) 321, https://doi.org/10.1038/s41598-020-79682-4.

[106] F. Hu, et al., Generating novel compounds targeting SARS-CoV-2 main protease based on imbalanced dataset, 2020 IEEE Int. C. Bioinform. (2020) 432–436, https://doi.org/10.1109/BIBM49941.2020.9313317.

[107] K. Gao, et al., Generative network complex for the automated generation of drug-like molecules, J. Chem. Inf. Model. 60 (12) (2020) 5682–5698, https://doi.org/10.1021/acs.jcim.0c00599.

[108] Y. Li, et al., Learning to Design Drug-like Molecules in Three-Dimensional Space Using Deep Generative Models, 2021, https://doi.org/10.48550/arXiv.2104.08474 *arXiv preprint arXiv:2104.08474*.
[109] Y. Li, et al., Structure-based de novo drug design using 3D deep generative models, Chem. Sci. 12 (2021) 13664–13675, https://doi.org/10.1039/D1SC04444C.
[110] J. Arús-Pous, et al., Randomized SMILES strings improve the quality of molecular generative models, J. Cheminf. 11 (1) (2019) 1–13, https://doi.org/10.1186/s13321-019-0393-0.
[111] L. David, et al., Applications of deep-learning in exploiting large-scale and heterogeneous compound data in industrial pharmaceutical research, Front. Pharmacol. 10 (2019) 1303, https://doi.org/10.3389/fphar.2019.01303.
[112] Y. Li, et al., Multi-objective de novo drug design with conditional graph generative model, J. Cheminf. 10 (1) (2018) 1–24, https://doi.org/10.1186/s13321-019-0393-0.
[113] G. Chen, et al., Machine-learning-assisted de novo design of organic molecules and polymers: opportunities and challenges, Polymers 12 (1) (2020) 163, https://doi.org/10.3390/polym12010163.
[114] M. Staszak, et al., Machine learning in drug design: use of artificial intelligence to explore the chemical structure–biological activity relationship, Wires. Comput. Mol. Sci. (2021) e1568, https://doi.org/10.1002/wcms.1568.
[115] V. Gallego, et al., AI in drug development: a multidisciplinary perspective, Mol. Divers. 25 (3) (2021) 1461–1479, https://doi.org/10.1007/s11030-021-10266-8.
[116] E. Lin, et al., Relevant applications of generative adversarial networks in drug design and discovery: molecular de novo design, dimensionality reduction, and de novo peptide and protein design, Molecules 25 (14) (2020) 3250, https://doi.org/10.3390/molecules25143250.
[117] P. Renz, et al., On failure modes in molecule generation and optimization, Drug Discov. Today 32 (2019) 55–63, https://doi.org/10.1016/j.ddtec.2020.09.003.
[118] H. Chen, et al., The rise of deep learning in drug discovery, Drug Discov. Today 23 (6) (2018) 1241–1250, https://doi.org/10.1016/j.drudis.2018.01.039.
[119] S. Vatansever, et al., Artificial intelligence and machine learning-aided drug discovery in central nervous system diseases: state-of-the-arts and future directions, Med. Res. Rev. 41 (3) (2021) 1427–1473, https://doi.org/10.1002/med.21764.
[120] S. Raschka, B. Kaufman, Machine learning and AI-based approaches for bioactive ligand discovery and GPCR-ligand recognition, Methods 180 (2020) 89–110, https://doi.org/10.1016/j.ymeth.2020.06.016.
[121] Y.-C. Lo, et al., Machine learning in chemoinformatics and drug discovery, Drug Discov. Today 23 (8) (2018) 1538–1546, https://doi.org/10.1016/j.drudis.2018.05.010.
[122] L. de Souza Neto, et al., Silico strategies to support fragment-to-lead optimization in drug discovery, Front. Chem. 8 (2020) 93, https://doi.org/10.3389/fchem.2020.00093.
[123] R. Lopez, et al., Enhancing scientific discoveries in molecular biology with deep generative models, Mol. Syst. Biol. 16 (9) (2020), e9198, https://doi.org/10.15252/msb.20199198.
[124] R.S.K. Vijayan, et al., Enhancing Preclinical Drug Discovery with Artificial Intelligence, Drug Discov. Today, 2021, https://doi.org/10.1016/j.drudis.2021.11.023.
[125] M.K. Tripathi, et al., Evolving scenario of big data and Artificial Intelligence (AI) in drug discovery, Mol. Divers. (2021) 1–22, https://doi.org/10.1007/s11030-021-10256-w.
[126] T.S. Toh, et al., Looking beyond the hype: applied AI and machine learning in translational medicine, EBioMedicine 47 (2019) 607–615, 10.1016/j.ebiom.2019.08.027.
[127] Y. Chen, J. Kirchmair, Cheminformatics in natural product-based drug discovery, Mol. Inform. 39 (12) (2020) 2000171, https://doi.org/10.1002/minf.202000109.
[128] M. Batool, et al., A structure-based drug discovery paradigm, Int. J. Mol. Sci. 20 (11) (2019) 2783, https://doi.org/10.3390/ijms20112783.
[129] W. Gao, C.W. Coley, The synthesizability of molecules proposed by generative models, J. Chem. Inf. Model. 60 (12) (2020) 5714–5723, https://doi.org/10.1021/acs.jcim.0c00174.
[130] G.M. Ghiandoni, et al., RENATE: a pseudo-retrosynthetic tool for synthetically accessible de Novo design, Mol. Inform. (2021) 2100207, https://doi.org/10.1002/minf.202100207.
[131] C.H. Liu, et al., Retrognn: Approximating retrosynthesis by graph neural networks for de novo drug design, 2020 *arXiv preprint arXiv:2011.13042*.
[132] S. Dara, et al., Machine learning in drug discovery: a review, Artif. Intell. Rev. (2021) 1–53, https://doi.org/10.1007/s10462-021-10058-4.
[133] G.M. Ghiandoni, et al., Enhancing reaction-based de novo design using a multi-label reaction class recommender, J. Comput. Aided Mol. Des. 34 (7) (2020) 783–803, https://doi.org/10.1007/s10822-020-00300-6.
[134] L.H. Mervin, et al., Uncertainty quantification in drug design, *Drug Discov. Today* 26 (2) (2021) 474–489, https://doi.org/10.1016/j.drudis.2020.11.027.
[135] N. Brown, et al., GuacaMol: benchmarking models for de novo molecular design, J. Chem. Inf. Model. 59 (3) (2019) 1096–1108, https://doi.org/10.1021/acs.jcim.8b00839.
[136] D. Polykovskiy, et al., Molecular sets (MOSES): a benchmarking platform for molecular generation models, Front. Pharmacol. 11 (2020) 1931, https://doi.org/10.3389/fphar.2020.565644.
[137] J. Meyers, et al., De novo molecular design and generative models, Drug Discov. Today 26 (11) (2021) 2707–2715, https://doi.org/10.1016/j.drudis.2021.05.019.
[138] K. Preuer, et al., Fréchet ChemNet distance: a metric for generative models for molecules in drug discovery, J. Chem. Inf. Model. 58 (9) (2018) 1736–1741, https://doi.org/10.1021/acs.jcim.8b00234.
[139] T. Cieplinski, et al., We Should at Least Be Able to Design Molecules that Dock Well, 2020 *arXiv preprint arXiv:2006.16955*.
[140] X. Liu, et al., DrugEx v2: de novo design of drug molecule by pareto-based multi-objective reinforcement learning in polypharmacology, J. Cheminf. 13 (2021) 85, https://doi.org/10.1186/s13321-021-00561-9.
[141] V.D. Mouchlis, et al., Advances in de novo drug design: from conventional to machine learning methods, Int. J. Mol. Sci. 22 (4) (2021) 1676, https://doi.org/10.3390/ijms22041676.
[142] P. Jia, et al., The roles of computer-aided drug synthesis in drug development, Green Synth. Catal. (2021), https://doi.org/10.1016/j.gresc.2021.11.007.
[143] J. Yasonik, Multiobjective de novo drug design with recurrent neural networks and nondominated sorting, J. Cheminf. 12 (1) (2020) 1–9, https://doi.org/10.1186/s13321-020-00419-6.