Available online at www.sciencedirect.com

ScienceDirect

Current Opinion in Biotechnology

Generative artificial intelligence for small molecule drug design

Ganesh Chandan Kanakala*, Sriram Devata*, Prathit Chatterjee and Udaykumar Deva Priyakumar

In recent years, the rapid advancement of generative artificial intelligence (GenAI) has revolutionized the landscape of drug design, offering innovative solutions to potentially expedite the discovery of novel therapeutics. GenAI encompasses algorithms and models that autonomously create new data, including text, images, and molecules, often mirroring characteristics of existing datasets. This comprehensive review delves into the realm of GenAI for drug design, emphasizing recent advancements and methodologies that have propelled the field forward. Specifically, we focus on three prominent paradigms: transformers, diffusion models, and reinforcement learning algorithms, which have been exceptionally impactful in the last few years. By synthesizing insights from a myriad of studies and developments, we elucidate the potential of these approaches in accelerating the drug discovery process. Through a detailed analysis, we explore the current state and future directions of GenAI in the context of drug design, highlighting its transformative impact on pharmaceutical research and development.

Address
Centre for Computational Natural Sciences and Bioinformatics, International Institute of Information Technology, Hyderabad 500032, Telangana, India

Corresponding author: U. Deva Priyakumar (deva@iiit.ac.in)
* These authors contributed equally.

Current Opinion in Biotechnology 2024, **89**:103175

This review comes from a themed issue on **Pharmaceutical Biotechnology**

Edited by **Michael Krogh Jensen**

Available online xxxx

https://doi.org/10.1016/j.copbio.2024.103175

Introduction

Brief overview

Machine learning is a subset of artificial intelligence (AI) focused on creating algorithms to learn from data for predictions or decisions without explicit programming of predefined sets of rules. AI encompasses a broader range of technologies for tasks requiring human-like intelligence, like understanding language or recognizing patterns. The first AI age (1960–1990) emphasized search algorithms, optimization, and 'Good-Old Fashioned AI' (GOFAI), driven by mathematical principles and rule-based systems, with risks primarily in software quality [1–3]. The second age (1980–2000) transitioned to machine learning, relying on domain knowledge and testing, but faced challenges of incomplete data and biases [3]. In addition to the new dataset generation, the ongoing third age marked an advancement in algorithms with deep learning and generative AI (GenAI). It benefits from vast data and computational power but introduces new risks from uncurated big data, intellectual property issues, and misuse concerns, emphasizing the need to understand and mitigate these risks [3].

What is generative artificial intelligence?

GenAI, a short form of generative artificial intelligence, is a subset of AI that creates new data resembling a given dataset instead of merely recognizing patterns or making predictions. Furthermore, it diverges from traditional AI by also creating entirely new data instances resembling the training data. These models learn the data's underlying structure to produce novel, synthetic examples rather than direct copies, accelerating innovation across various domains. Techniques like generative adversarial networks (GANs) [4], variational autoencoders (VAEs) [5], transformers, and reinforcement learning (RL) have diverse applications, including drug discovery. GenAI revolutionizes drug discovery by leveraging machine learning to generate novel molecular structures. Primarily, we identify potential data sources in 1D, 2D, or 3D representations to create a curated dataset. Thereafter, we train any generative model, such as generative pretrained transformers (GPTs), recurrent neural networks (RNNs), VAEs, graph neural networks (GNNs), GANs, deep belief networks, equivariant GNNs, Markov Chain Monte-Carlo models, probabilistic graphical models, and so on, for molecule generation, preferably with a way to condition them on desired properties. Figure 1 summarizes the typical approach for molecule generation tasks using GenAI methods.

Generative artificial intelligence in drug design

Representations of molecules

Advances in deep learning architectures from other domains can be used for generating drug molecules and predicting their properties by using various representations of molecules. Text-based representations allow for the domain of drug molecules to benefit from advances in natural language

Figure 1

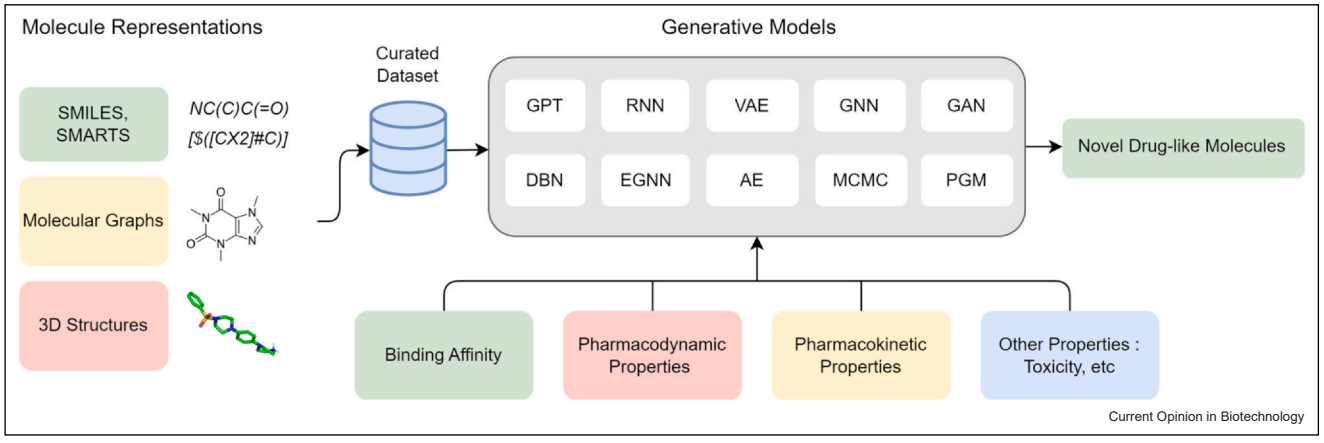

Overview of GenAI for drug design.

processing. Graph-based representations of drug molecules facilitate the use of GNNs, inspired by advances in social and citation networks as well as scene graphs. Tools like RDKit aid in converting molecular data between different representations. While early works utilized molecular fingerprints, recent literature predominantly favors more comprehensive representations that contain more information about the molecules. The subsequent sections detail these representations and their respective chemical information content.

Simplified Molecular-Input Line-Entry System and Self-Referencing Embedded Strings (1D/string)

Simplified Molecular-Input Line-Entry System (SMILES) [6] represents a molecule's structure as a string while encapsulating atoms, bonds, and stereochemistry. Although each molecule can have multiple SMILES strings, canonicalization algorithms yield a single representative canonical SMILES string. However, not all SMILES are valid due to syntax discrepancies. SELFIES [7], on the other hand, ensures every string in the correct syntax corresponds to a valid molecule, offering a robust alternative.

2D molecular graphs (2D)

A 2D molecular graph comprises nodes N representing atoms and edges E representing bonds. It typically involves three matrices: an adjacency matrix indicating vertex adjacency; a node matrix storing node attributes like atomic number, valency, and formal charge; and an edge matrix containing bond properties like bond type.

3D molecular graphs (3D)

In addition to depicting atom connectivity and topology, 3D graphs also incorporate spatial positions of atoms. This spatial data can be represented explicitly with atom coordinates in the node matrix or through relative 3D information like bond angles and lengths, and dihedral angles. While obtaining 3D geometric information requires expensive computations and experiments, it proves valuable in molecular analysis tasks due to its inclusion of ground-state geometries with lowest energies.

Databases

There exist multiple databases containing molecules meeting specific criteria (e.g. all molecules that have < 10 heavy atoms) or serving various functionalities. Understanding the choice of training database is crucial for GenAI models, which learn from underlying data distributions to generate new samples. While pharmaceutical companies possess proprietary data, publicly available databases offer options for (pre)training these models to generate potential new drug molecules. The GDB-17 database enumerates 166 billion organic small molecules of up to 17 atoms (C, N, O, S, halogens) [8], while QM-9 is a subset with about 140 000 molecules up to 9 atoms (C, N, O, F) [9]. These databases encompass molecules within their defined chemical space and may include quantum chemical properties. Large databases like GDB-17 aid in pre-training GenAI models to grasp underlying molecular patterns, which can be fine-tuned for property-specific molecule generation. The ZINC database [10] and its variants offer multi-billion-scale collections of commercially available drug-like molecules and include vendor information for obtaining the molecules. MOSES database [11] comprises around 2 million potential hit compounds suitable for further Chemical absorption, distribution, metabolism, excretion, and toxicity optimization, while ChEMBL [12] is a curated repository of bioactive molecules, with about 2.4 million entries and experimental bioactivity data. PubChem's Compound database [13] boasts over 111 million unique chemical structures alongside bioactivity data from assays. DrugBank [14] houses over 10 000 approved and investigational drugs and 1.4 million drug–drug interactions.

Table 1

Summary of various databases available for training and testing generative models for small (drug) molecules.

Name of Database	Number of Relevant Entries	Comments
GDB-17 [9]	166 B	Organic small molecules of up to 17 atoms; can be used for pre-training
ZINC [10]	37 B	Commercially available compounds; Contains database search tools
PubChem [13]	111 M	Contains bioactivity measurements; Has database search tools
CrossDocked [15]	22.5 M	Contains poses of ligands bound to binding targets
ChEMBL [12]	2.4 M	Drug-like bioactive molecules; Contains bioactivity measurements
MOSES [11]	2 M	Potential hit compounds; Opportunities for further ADMET optimizations
QM-9 [9]	140 K	Molecules of up to 9 atoms; also contains quantum chemical properties
DrugBank [14]	10 K	Approved and investigational drugs; Also contains 1.4 M drug-drug interactions

ADMET, Chemical absorption, distribution, metabolism, excretion, and toxicity.

In addition to databases containing drug-related information, the CrossDocked database [15] offers 22.5 million ligand poses docked into Protein Data Bank [16] binding pockets, useful for generating molecules targeting specific pockets. While commercial collections like Enamine REAL Space [17], CAS [18], and Beilstein archives [19] are not detailed due to limited accessibility, Table 1 summarizes widely used open databases for training generative models in molecule design.

Generative methods for drug design

Generative models play a vital role in drug design, particularly in *de novo* drug design. Early approaches using RNNs effectively generated 1D sequences like SMILES but struggled with larger molecules. Therefore, long short-term memory models and other variants of RNNs became a popular choice of sequence-based models due to their capabilities for modeling longer sequences. VAEs [5] offer an alternative approach by mapping molecular structures to a low-dimensional latent space, facilitating the generation of novel molecules by sampling the diverse latent space. Various VAE variants have been proposed to overcome limitations and improve molecule generation. GVAEs [20] proved to be a significant step in improving the validity of the molecules by incorporating a parse-tree representation for molecules. Constrained graph VAEs (CGVAEs) take a step forward by combining VAEs and graph generation using Gated Graph Neural Nets [21] for novel molecular graph generation. Similarly, junction tree VAEs innovate molecular VAEs by using a junction tree (a tree-structured scaffold) representation of a molecule followed by a message-passing network for molecule generation. Deep learning-based inorganic material generator [22] extends this by using CGVAEs for conditional inorganic molecule generation.

GANs [4], popular for image generation, consist of a generator network and a discriminator network. The generator creates new data, while the discriminator identifies real data from the fake ones. This adversarial training aims to fool the discriminator. Early applications in molecule generation involved using an adversarial autoencoder to generate molecular fingerprints [23], which were used to screen compounds from PubChem. Because of inherent similarities between GANs and actor–critic architectures, researchers have explored combining them, giving rise to innovations such as optimizing GANs using RL. ORGAN [24] is one of the earliest methods to do this, which employed a GAN for generation and RL for biasing the molecular generation towards desired metrics like solubility, drug likeliness, and synthesizability. This success has led to numerous GAN variants showcasing their potential in drug design. MolGAN [25] is another example that uses a GAN for molecule graph generation within an RL setting for the same properties optimization. LatentGAN [26] combines an autoencoder with a GAN for SMILES generation for specific targets such as Epidermal Growth Factor Receptor (EGFR), 5-hydroxytryptamine (serotonin) receptor 1A, and Sphingosine-1-phosphate receptor 1. Mol-CycleGAN [27], for instance, can generate highly similar molecules with desired properties given a target molecule. One of the main challenges with training GANs is the phenomenon of mode collapse when the generator keeps generating a small subset of molecules and continues tricking the discriminator and getting stuck in local minima. These issues can be addressed by using concepts from genetic algorithms and adaptive training data [28].

Flow-based models are a relatively recent development in molecule generation. In a way, they are similar to VAEs in learning the representation; however, the main difference is that these models leverage normalizing flows [29] to learn a bijective mapping between data and a simple distribution. Flow-based models construct a series of invertible transformations that map a simple distribution to the data distribution. Flow-based models were successful in the image domain. To extend these to molecules, Madhawa et al. proposed GraphNVP [30], using an adjacency matrix and node feature matrix as the data to be generated using invertible neural networks, which can be decoded to molecular graphs. Several pieces of the literature suggest that flow-based methods have been successfully used for molecular design for different modalities, such as graphs and sequences. MoFlows [31] uses a flow-based model to generate bonds, then atoms, and post hoc chemical validity checks for guaranteed valid molecular graph generation.

FastFlows [32] extends these normalizing flows to sequence generation by using SELFIES, making it extremely efficient for generating molecules in bulk. MolGrow [33] innovates by using a hierarchical approach to molecular graph generation recursively using normalizing flows by splitting a node in each iteration, maintaining chemical validity and structural diversity.

Recent advances in molecule generation go beyond the methods mentioned earlier. GPTs, diffusion models, and RL are now showing promise. The subsequent sections will delve into a comprehensive discussion of these methodologies in the context of molecular generation.

Transformers for drug design

Transformers [34] emerged as a powerful solution to overcome the limitations with RNN-based models with their strength in capturing long-range dependencies in sequences.

Transformers feature an encoder–decoder architecture, the latter, as shown in the Figure 2a, used for sequence generation. A typical approach to molecule generation using transformers includes (1) Encoding the molecules and property information followed by positional encodings, (2) Input to the transformer decoder with several multihead attention modules followed by a feed-forward neural network for next token classification. Multiple self-attention 'heads' allow for learning distinct molecule aspects, enhancing information extraction. This architecture offers versatility, with the encoder used for representation learning and downstream tasks like property predictions [35], while the decoder is employed for tasks such as molecule generation [36].

Transformers have been used widely for conditional molecule generation and unconditional molecule generation.

- **Unconditional molecule generation** aims to create diverse molecules without specific properties, maximizing validity, novelty, and exploration of chemical space. It serves as a valuable pre-training step for conditional molecule generation.
- **Conditional molecule generation** tailors generative models to prioritize molecules with desired properties by incorporating property information [36]. Simple methods embed properties as starting tokens, as illustrated in Figure 2a. This can be achieved by embedding properties as starting tokens or improved by encoding graph-based representations (GraphGPT) [37] into the sequence model. Further innovations include methods like cMolGPT [38], which injects conditional information directly into the decoder's attention layer and approaches like CMGN [39], which employ large-scale pretraining and fine-tuning for target-specific molecule generation. Various methods for incorporating conditional information define the range of generative tasks addressed by molecular GPTs.

Figure 2

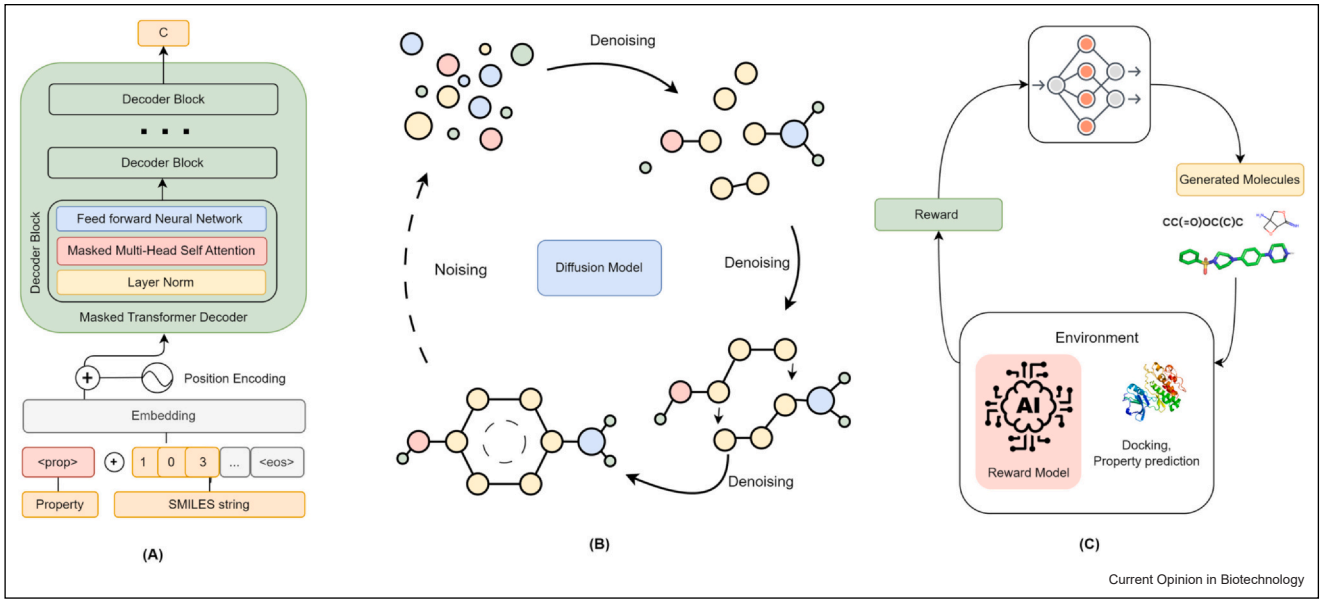

Recent advancements in generative models for drug design. **(a)** GPT-based molecular design. **(b)** Diffusion-based molecular design. **(c)** RL for molecular design.

Diffusion models for drug design

Diffusion models [40], parameterized Markov chains trained via variational inference, excel at generating novel data points within a specific domain. During training, a forward diffusion process progressively injects Gaussian noise into real data, ultimately transforming it into an isotropic Gaussian distribution. The subsequent reverse process trains the model to iteratively remove this added noise, essentially performing denoising at each step. This learned denoising ability allows the model to take random noise as input and, through a series of denoising steps, generate new data points that statistically resemble the training data.

Diffusion models, originally successful in computer vision, have been swiftly adapted to 3D molecule generation. E(3) equivariant diffusion models pioneered this adaptation [41], achieving state-of-the-art (SOTA) 3D conditional molecule generation on QM9 dataset. Figure 2b showcases a general application of diffusion methods to the molecular domain. As shown in the figure, the diffusion model is trained to denoise a molecule periodically from a noised variant with noisy coordinates and atom types, perturbing the atom coordinates and types until it becomes a valid molecule, thereby translating the concept of noising and denoising image pixels to atomic coordinates and types. This success has sparked further advancements, with newer models like latent diffusion models (LDMs) [42] being used for molecule generation as shown by geometric LDM [43], which use point structured latent space and autoencoders operating there to resolve large molecules, and geometry complete diffusion models [43] incorporating VAEs and SE(3)-Equivariant structures to enhance conditioning capabilities. Molecular diffusion models (MDMs) were further introduced to address challenges faced by diffusion models in handling large molecules and diversity [44]. MDMs utilize equivariant encoders to encode interatomic relations and incorporate a latent variable for controlling representations in each diffusion/reverse step, ensuring the exploration of diverse 3D molecule geometries. Their approach achieved SOTA performance on QM9 and Geom-Drugs datasets. Diffusion models have further expanded into target-specific molecule generation in 3D. DiffSBDD and TargetDiff are SE(3)-Equivariant models for generating ligands conditioned on protein pockets [45,46], while DiffDOCK frames docking as a generative modeling problem focusing on ligand poses [47]. Additionally, torsional diffusion for conformation generation [48] and steering diffusion model training with physical and statistical information [49] were further introduced. With their versatility and potential, diffusion models are poised to play a pivotal role in future molecule discovery and design.

Reinforcement learning for drug design

RL is a learning paradigm where an *agent* (also referred to as an *actor* or *decision maker*) learns *policies* to maximize cumulative rewards from the environment (also referred as *critic*) for a specific task. As shown in Figure 2c, the agent selects actions based on its policy, and the environment provides rewards and new states accordingly. The agent updates its policy using intermediate and final rewards. The environment's design includes state representation, action definitions, transition dynamics, and reward function, while the agent's design involves choosing the policy network architecture for action selection. In molecule generation, RL often uses molecular graphs or string-based representations, where actions include adding or removing edges in graphs or concatenating tokens in strings. RL is advantageous for sequential decision problems in molecular design, as each decision moves partial molecules closer to desired properties. Chemistry-aware RL environments can incorporate constraints on nondifferentiable properties (molecular structure and valency) in the reward function. Unlike other methods, RL allows exploration beyond training data, provided the problem formulation and reward function support it. ReLeaSE [50] is one of the earliest works in RL for a molecular design where the agent generates SMILES strings, and a predictive model in the environment predicts a property value that serves as a reward for the agent to generate compounds with the desired properties. Other early works like Graph Convolutional Policy Network [51] employed graph convolutional networks to generate molecular graphs with specific properties like logP and drug-likeness. With the reward functions becoming more realistic and computationally heavy, recent works introduce an RL-based active learning system for hit molecule generation [52,53] to reduce the evaluations needed in the generation process. MoleGuLAR [54] adopts a multiobjective approach by periodically changing reward properties. DeepSPInN [55] uses Monte Carlo Tree Search (MCTS) for candidate molecule generation from spectral data. FREED and FREED++ combine molecular fragments to generate high-affinity molecules [56,57]. MolDQN and iterative RL utilize Q-learning for molecule generation and optimization [58,59]. Apart from generating molecules, RL has also been used for other tasks, such as predicting structures of metal clusters [60] and for molecular geometry optimization [61].

With transformers receiving a lot of attention in recent years, we also highlight works that use RL along with transformers. ChemRLformer [62] uses RL to train a transformer to generate string representations of ligands that have high binding affinities to multiple docking targets. Molecular design using Reinforcement Learning with Multiple GPT agents [63] uses multiple generative pre-trained transformer agents to generate molecules optimized for particular properties.

Conclusion

The integration of GenAI into drug discovery promises accelerated innovation in pharmaceutical research, leveraging diverse molecule representations like SMILES, SELFIES, and graphs. Comprehensive databases fuel GenAI by providing vast training data, enhancing the reliability of

generated compounds. Generative methods, including transformers, diffusion, and RL, revolutionize lead identification. Despite progress, challenges persist, requiring addressing reliability, interpretability, data security, and ethical considerations for responsible GenAI application in drug discovery. Nonetheless, the potential for GenAI to expedite medication development remains compelling, heralding a future of faster, more precise drug discovery.

Author contribution

UDP formulated the review and its flow. GCK, SD, and PC wrote the manuscript, and UDP supervised it. All authors reviewed and approved the final version.

Data Availability

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

Declaration of Competing Interest

The authors declare no competing interests.

Declaration of Generative AI and AI-assisted technologies in the writing process

The authors declare using ChatGPT (https://chat.openai.com/), for improved readability of the text and making it concise. After using the service, the authors have reviewed and edited the text, and take full responsibility for the content of the publication.

Acknowledgements

The authors thank DST-SERB, India (CRG/2021/008036) and IHub-Data, IIIT Hyderabad for financial support.

References and recommended reading

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

- • of special interest
- •• of outstanding interest

1. Boden MA: **GOFAI**. Cambridge University Press; 2014:89-107.
2. Fradkov AL: **Early history of machine learning**. *IFAC PapersOnLine* 2020, **53**:1385-1390.
3. Cao L: **A new age of AI: features and futures**. *IEEE Intell Syst* 2022, **37**:25-37.
4. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y: **Generative adversarial networks**. *Commun ACM* 2020, **63**:139-144.
5. Kingma DP, Welling M: **Auto-encoding Variational Bayes**. *aRxiv* 2022, pre-print: not peer reviewed. https://arxiv.org/abs/1312.6114v11.
6. Weininger D: **Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules**. *J Chem Inf Comput Sci* 1988, **28**:31-36.
7. Krenn M, Häse F, Nigam A, Friederich P, Aspuru-Guzik A: **Self-referencing embedded strings (selfies): a 100% robust molecular string representation**. *Mach learn Sci Tech* 2020, **1**:045024.
8. Ruddigkeit L, van Deursen R, Blum LC, Reymond J-L: **Enumeration of 166 billion organic small molecules in the chemical universe database gdb-17**. *J Chem Inf Model* 2012, **52**:2864-2875.
9. Ramakrishnan R, Dral PO, Rupp M, vonLilienfeld OA: **Quantum chemistry structures and properties of 134 kilo molecules**. *Sci Data* 2014, **1**:140022.
10. Irwin JJ, Shoichet BK: **ZINC — a free database of commercially available compounds for virtual screening**. *J Chem Inf Model* 2005, **45**:177-182.
11. Polykovskiy D, Zhebrak A, Sanchez-Lengeling B, Golovanov S, Tatanov O, Belyaev S, Kurbanov R, Artamonov A, Aladinskiy V, Veselov M, *et al*.: **Molecular sets (moses): a benchmarking platform for molecular generation models**. *Front Pharmacol* 2020, **11**:565644.
12. Zdrazil B, Felix E, Hunter F, Manners EJ, Blackshaw J, Corbett S, de Veij M, Ioannidis H, Lopez DM, Mosquera JF, Magarinos MP, Bosc N, Arcila R, Kizilören T, Gaulton A, Bento AP, Adasme MF, Monecke P, Landrum GA, Leach AR: **The ChEMBL Database in 2023: a drug discovery platform spanning multiple bioactivity data types and time periods**. *Nucleic Acids Res* 2023, **52**:D1180-D1192.
13. Kim S, Chen J, Cheng T, Gindulyte A, He J, He S, Li Q, Shoemaker BA, Thiessen PA, Yu B, Zaslavsky L, Zhang J, Bolton EE: **PubChem in 2021: new data content and improved web interfaces**. *Nucleic Acids Res* 2020, **49**:D1388-D1395.
14. Knox C, Wilson M, Klinger CM, Franklin M, Oler E, Wilson A, Pon A, Cox J, Chin NEL, Strawbridge SA, Garcia-Patino M, Kruger R, Sivakumaran A, Sanford S, Doshi R, Khetarpal N, Fatokun O, Doucet D, Zubkowski A, Rayat DY, Jackson H, Harford K, Anjum A, Zakir M, Wang F, Tian S, Lee B, Liigand J, Peters H, Wang RQR, Nguyen T, So D, Sharp M, da Silva R, Gabriel C, Scantlebury J, Jasinski M, Ackerman D, Jewison T, Sajed T, Gautam V, Wishart DS: **DrugBank 6.0: the DrugBank Knowledgebase for 2024**. *Nucleic Acids Res* 2023, **52**:D1265-D1275.
15. Francoeur PG, Masuda T, Sunseri J, Jia A, Iovanisci RB, Snyder I, Koes DR: **Three-dimensional convolutional neural networks and a cross-docked data set for structure-based drug design**. *J Chem Inf Model* 2020, **60**:4200-4215.
16. Berman HM, Westbrook J, Feng Z, Gilliland G, Bhat TN, Weissig H, Shindyalov IN, Bourne PE: **The protein data bank**. *Nucleic Acids Res* 2000, **28**:235-242.
17. Enamine Real Space and Real Database**.** https://enamine.net/compound-collections/real-compounds, 2024 (Accessed 21-03-2024).
18. **Cas registry system**. *J Chem Inf Comput Sci* (1) 1978, **18**:58 https://pubs.acs.org/doi/pdf/10.1021/ci60013a609.
19. Heller SR: **The Beilstein Online Database**. American Chemical Society; 1990.
20. Kusner MJ, Paige B, Hernández-Lobato JM: **Grammar variational autoencoder**. *aRxiv* 2017, pre-print: not peer reviewed. https://arxiv.org/abs/1703.01925v1.
21. Li Y, Tarlow D, Brockschmidt M, Zemel RS: **Gated graph sequence neural networks.** In Proceedings of the 4th International Conference on Learning Representations, ICLR 2016 May 2–4; San Juan, Puerto Rico, Conference Track Proceedings Edited by Bengio Y, LCun Y; 2016. ⟨http://arxiv.org/abs/1511.05493⟩.
22. Pathak Y, Juneja KS, Varma G, Ehara M, Priyakumar UD: **Deep learning enabled inorganic material generator**. *Phys Chem Chem Phys* 2020, **22**:26935-26943.
23. Kadurin A, Aliper A, Kazennov A, Mamoshina P, Vanhaelen Q, Khrabrov K, Zhavoronkov A: **The cornucopia of meaningful leads: applying deep adversarial autoencoders for new molecule development in oncology**. *Oncotarget* 2017, **8**:10883-10890.
24. Guimaraes GL, Sanchez-Lengeling B, Outeiral C, Farias PLC, Aspuru-Guzik A: **Objective-reinforced generative adversarial networks (organ) for sequence generation models**. *aRxiv* 2017, pre-print: not peer reviewed. https://arxiv.org/abs/1705.10843v3.
25. Cao MD, Kipf T: **MolGAN: An implicit generative model for small molecular graphs**. *aRxiv* 2022, pre-print: not peer reviewed. https://arxiv.org/abs/1805.11973v2.

26. Prykhodko O, Johansson SV, Kotsias P-C, Arús-Pous J, Bjerrum EJ, Engkvist O, Chen H: **A de novo molecular generation method using latent vector based generative adversarial network**. *J Cheminform* 2019, **11**:74.

27. Maziarka Ł, Pocha A, Kaczmarczyk J, Rataj K, Danel T, Warchoł M: **Mol-CycleGAN: a generative model for molecular optimization**. *J Cheminform* 2020, **12**:2, https://doi.org/10.1186/s13321-019-0404-1

28. Blanchard AE, Stanley C, Bhowmik D: **Using GANs with adaptive training data to search for new molecules**. *J Cheminform* 2021, **13**:14, https://doi.org/10.1186/s13321-021-00494-3

29. Rezende DJ, Mohamed S: **Variational inference with normalizing flows**. *aRxiv* 2016, pre-print: not peer reviewed. https://arxiv.org/abs/1505.05770v6.

30. Madhawa K, Ishiguro K, Nakago K, Abe M: **GraphNVP: An invertible flow model for generating molecular graphs**. *aRxiv* 2019, pre-print: not peer reviiewed. https://arxiv.org/abs/1905.11600v1.

31. Zang C, Wang F: **Moflow: An invertible flow model for generating molecular graphs.** In: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM; 2020.

32. Yu J, Zheng Y, Wang X, Li W, Wu Y, Zhao R, Wu L: **Fastflow: Unsupervised anomaly detection and localization Via 2d normalizing flows**. *aRxiv* 2021, pre-print: not peer reviewed https://arxiv.org/abs/2111.07677v2.

33. Kuznetsov M, Polykovskiy D: **Molgrow: A graph normalizing flow for hierarchical molecular generation**. *Proc AAAI Conf Artif Intell* 2021, **35**:8226-8234 ⟨https://ojs.aaai.org/index.php/AAAI/article/view/17001⟩.

34. Vaswani A, Shazeer N, Parmar P, Uszkoreit J, Jones L, Gomez AN, Kaiser L, Polosukhin I: **Attention is all you need**. *aRxiv* 2023,, https://arxiv.org/abs/1706.03762v7.

35. Chithrananda S, Grand G, Ramsundar B: **ChemBERTa: large-scale self-supervised pretraining for molecular property prediction**. *aRxiv* 2020, pre-print: not peer-reviewed. https://arxiv.org/abs/2010.09885v2.

36. • Bagal V, Aggarwal R, Vinod PK, Priyakumar UD: **Molgpt: molecular generation using a transformer-decoder model**. *J Chem Inf Model* 2022, **62**:2064-2076.

MolGPT employs a transformer decoder architecture for multi-property conditional molecule generation using scaffold input and properties of interest such as LogP, QED, TPSA, and SAS. The results show that MolGPT can generate valid and novel molecules on Guacamol and MOSES datasets for various properties and scaffolds.

37. Lu H, Wei Z, Wang X, Zhang K, Liu H: **Graphgpt: a graph enhanced generative pretrained transformer for conditioned molecular generation**. *Int J Mol Sci* 2023, **24**:16761.

38. •• Wang Y, Zhao H, Sciabola S, Wang W: **cMolGPT: a conditional generative pre-trained transformer for target-specific de novo molecular generation**. *Molecules* 2023, **28**:4430.

cMolGPT is a transformed-based network that implements keys and values for multihead attention modules using target-specific information. cMolgGPT is capable of generating SMILES strings that correspond to drug-like active compounds. Results suggest that this method can generate active drug-like molecules for different targets such as Epidermal Growth Factor Receptor, HTR1A, and S1PR1 and can be a valuable tool for *de novo* drug design.

39. • Yang M, Sun H, Liu X, Xue X, Deng Y, Wang X: **CMGN: a conditional molecular generation net to design target-specific molecules with desired properties**. *Brief Bioinforma* 2023, **24**:bbad185.

CMGN is a conditional molecule generation method that uses a bidirectional transformer network for target specific conditional generation. It takes advantage of large scale pretraining to learn SMILES chemistry and later finetuned for target specificity. Additionally, CMGN is able to take multiple (one to three) fragments as input for fragment-based drug design and multi-objective lead optimization.

40. Ho J, Jain A, Abbeel P: **Denoising diffusion probabilistic models**. *aRxiv* 2020, pre-print: not peer reviewed. https://arxiv.org/abs/2006.11239v2.

41. Hoogeboom E, Satorras VG, Vignac C, Welling M: **Equivariant diffusion for molecule generation in 3D**. *aRxiv* 2022, pre-print: not peer reviewed. https://arxiv.org/abs/2203.17003v2.

42. Rombach R, Blattmann D, Lorenz D, Esser P, Ommer B: **High-resolution image synthesis with latent diffusion models**. *Proc IEEE/CVF Conf Comput Vis Pattern Recognit* 2022,10684-10695.

43. Xu M, Powers A, Dror R, Ermon S, Leskovec J: **Geometric latent diffusion models for 3D molecule generation**. *aRxiv* 2023, pre-print: not peer reviewed. https://arxiv.org/abs/2305.01140v1.

44. Huang L, Zhang T, Xu T, Wong KC: **MDM: Molecular diffusion model for 3D molecule generation**. *aRxiv* 2022, pre-print: not peer reviewed. https://arxiv.org/abs/2209.05710v1.

45. Schneuing A, Du Y, Harris C, Jamasb A, Igashov I, Du W, Blundell T, Lió P, Gomes C, Welling M, Bronstein M, Correia B: **Structure-based drug design with equivariant diffusion models**. *aRxiv* 2023, pre-print: not peer-reviewed. https://arxiv.org/abs/2210.13695v2.

46. •• Guan J, Qian WW, Peng X, Su Y, Peng J, Ma J: **3D Equivariant diffusion for target-aware molecule generation and affinity prediction**. *aRxiv* 2023,.

TargetDiff is a 3D equivariant diffusion model for target-aware molecule generation. The model learns a joint denoising of atom coordinates and categorical atom types using SE(3)-equivariant network. Results suggest that the proposed framework is able to generate realistic 3D molecules with higher binding affinities for the specified protein targets.

47. Corso G, Stärk H, Jing B, Barzilay R, Jaakkola T: **DiffDock: Diffusion steps, twists, and turns for molecular docking**. *aRxiv* 2023, pre-print: not peer reviewed. https://arxiv.org/abs/2210.01776v2.

48. Jing B, Corso G, Chang J, Barzilay R, Jaakkola T: **Torsional diffusion for molecular conformer generation**. *aRxiv* 2023, pre-print: not peer reviewed. https://arxiv.org/abs/2206.01729v2.

49. Wu L, Gong C, Liu X, Ye M, Liu Q: **Diffusion-based Molecule Generation with Informative Prior Bridges**. *aRxiv* 2022, pre-print: not peer reviewed. https://arxiv.org/abs/2209.00865v1.

50. Popova M, Isayev O, Tropsha A: **Deep reinforcement learning for de novo drug design**. *Sci Adv* 2018, **4**:eaap7885.

51. You J, Liu B, Ying R, Pande V, Leskovec J: **Graph convolutional policy network for goal-directed molecular graph generation**. *aRxiv* 2019, pre-print: not peer reviewed. https://arxiv.org/abs/1806.02473v3.

52. Dodds M, Guo J, Löhr T, Tibo A, Engkvist O, Janet JP: **Sample efficient reinforcement learning with active learning for molecular design**. *Chem Sci* 2024, **15**:4146-4160.

53. Viswanathan K, Goel M, Laghuvarapu S, Varma G, Priyakumar UD: **Streamlining pipeline efficiency: a novel model-agnostic technique for accelerating conditional generative and virtual screening pipelines**. *Sci Rep* 2023, **13**:21069.

54. Goel M, Raghunathan S, Laghuvarapu S, Priyakumar UD: **Molegular: molecule generation using reinforcement learning with alternating rewards**. *J Chem Inf Model* 2021, **61**:5815-5826.

55. • Devata S, Sridharan B, Mehta S, Pathak Y, Laghuvarapu S, Varma G, Priyakumar UD: **Deepspinn — deep reinforcement learning for molecular structure prediction from infrared and ^{13}c NMR spectra**. *Digit Disc* 2024, **3**:818-829.

DeepSPInN is an MCTS-based method that uses infrared and ^{13}c nmr spectra to generate candidate molecules. On the QM9 dataset, the authors achieve an accuracy of 91.5% with an average time of 77 seconds. This is the first study to not use any pre-existing spectral databases or molecular fragment knowledge bases for molecular structure elucidation.

56. Yang S, Hwang D, Lee S, Ryu S, Hwang SJ: **Hit and lead discovery with explorative RL and fragment-based molecule generation**. *Adv Neural Inf Process Syst* 2021, **34**:7924-7936.

57. •• Telepov A, Tsypin A, Khrabrov K, Yakukhnov S, Strashnov P, Zhilyaev P, Rumiantsev E, Ezhov D, Avetisian M, Popova O, Kadurin A: **Freed++: Improving RL agents for fragment-based molecule generation by thorough reproduction**. *aRxiv* 2024, pre-print: not peer reviewed. https://arxiv.org/abs/2401.09840v1.

FREED++ fixes the original FREED algorithm, increases the speed, and evaluates the improved version as well. This RL method attaches

multiple molecular fragments to generate molecules that have a high docking score to multiple binding targets.

58. Zhou Z, Kearnes S, Li L, Zare RN, Riley P: **Optimization of molecules via deep reinforcement learning**. *Sci Rep* 2019, **9**:10752.

59. Fang Y, Pan X, Shen H-B: **De novo drug design by iterative multiobjective deep reinforcement learning with graph-based molecular quality assessment**. *Bioinformatics* 2023, **39**:btad157.

60. Modee R, Verma A, Joshi K, Priyakumar UD: **Megen-generation of gallium metal clusters using reinforcement learning**. *Mach Learn Sci Tech* 2023, **4**:025032.

61. Modee R, Mehta S, Laghuvarapu S, Priyakumar UD: **Molopt: autonomous molecular geometry optimization using multiagent reinforcement learning**. *J Phys Chem B* 2023, **127**:10295-10303.

62. Ghugare R, Miret S, Hugessen A, Phielipp M, Berseth G: **Searching for high-value molecules using reinforcement learning and transformers**. *aRxiv* 2023, pre-print: not peer reviewed. https://arxiv.org/abs/2310.02902v1.

63. Hu X, Liu G, Zhao Y, Zhang H: **De novo drug design using reinforcement learning with multiple gpt agents**. *Adv Neural Inf Process Syst* 2024.