scientific reports

OPEN



Streamlining pipeline efficiency: a novel model-agnostic technique for accelerating conditional generative and virtual screening pipelines

Karthik Viswanathan, Manan Goel, Siddhartha Laghuvarapu, Girish Varma & U. Deva Priyakumar[⊠]

The discovery of potential therapeutic agents for life-threatening diseases has become a significant problem. There is a requirement for fast and accurate methods to identify drug-like molecules that can be used as potential candidates for novel targets. Existing techniques like high-throughput screening and virtual screening are time-consuming and inefficient. Traditional molecule generation pipelines are more efficient than virtual screening but use time-consuming docking software. Such docking functions can be emulated using Machine Learning models with comparable accuracy and faster execution times. However, we find that when pre-trained machine learning models are employed in generative pipelines as oracles, they suffer from model degradation in areas where data is scarce. In this study, we propose an active learning-based model that can be added as a supplement to enhanced molecule generation architectures. The proposed method uses uncertainty sampling on the molecules created by the generator model and dynamically learns as the generator samples molecules from different regions of the chemical space. The proposed framework can generate molecules with high binding affinity with ~a 70% improvement in runtime compared to the baseline model by labeling only ~30% of molecules compared to the baseline oracle.

Data-driven approaches have found profound success across multiple computer science domains—computer vision, natural language processing, signal processing, speech recognition, etc. These algorithms have made their way into drug discovery as well. With the increasing availability of open sources, well-curated datasets like ZINC¹, ChEMBL², and more have opened up avenues for using machine learning in different tasks like molecular property prediction, molecular structure prediction, retrosynthesis, and de novo molecule generation.

Conventionally, in order to identify potential molecules, high throughput screening (HTS) is performed on large databases in which every molecule in a database undergoes automated in vitro testing to find if they were a potential match³. High throughput screening is extremely expensive, inefficient, and has a low-hit rate⁴. This led researchers to formulate in silico computational approaches that could emulate protein-ligand interactions leading to high throughput screening experiments using physics-based property prediction methods to make the process more efficient⁵. Virtual screening is more cost-effective and efficient than high throughput screening. Further optimizations to virtual screening included clustering databases and using machine learning approaches to label molecules^{6,7}. Given the linear search time, finding a desirable molecule becomes extremely ineffective when virtual screening is performed on huge databases. Moreover, the lack of chemical diversity in datasets has made virtual screening non-universal⁸.

The fundamental idea behind HTS and HTVS is to exploit the molecules already known in the chemical space but this number is infinitesimal in comparison to the estimated size of the chemical space with about 10⁶⁰ synthesizable molecules⁹. Even the most exhaustive studies have been able to computationally evaluate only 10⁸ compounds¹⁰. Hence, the idea of de novo molecule generation emerged in which computational methods are used to generate molecules with certain properties. One way to achieve this is through genetic algorithms, which seek to evolve the ecosystem of pre-existing molecules into more desirable ones by introducing mutations in the

Center for Computational Natural Sciences and Bioinformatics, International Institute of Information Technology, Hyderabad 500032, India. [⊠]email: deva@iiit.ac.in

current generation^{11,12}. Unfortunately, genetic algorithms and their variants are prone to be stuck at local minima due to fixed initial populations and reverting mutations. To remove the initial requirement of a population set, deep generative models provide a vital improvement by forcing a non-linear relationship between molecular structures and properties¹³. In order to apply deep generative modeling to molecule generation, the two common representations are SMILES (simplified molecular-input line-entry system) strings and molecular graphs. SMILES strings possess their own grammar and semantics, and hence, this opens up the avenue for the application of natural language processing (NLP) based approaches like recurrent neural networks and transformers¹⁴. On the other hand, molecular graphs are generally heterogeneous graphs that can be used as an input to graph neural networks. Gupta et al.¹⁵ and Grisoni et al.¹⁶ used variants of recurrent neural networks to generate generic drug-like molecules while Bongini et al.¹⁷ and Mercado et al.¹⁸ used graph neural networks for molecule generation. However, drug molecules for a novel disease must possess a particular set of properties, and hence, methods were required to generate molecules with specified properties. This led to the application of more sophisticated generative models like variational autoencoders (VAE)¹⁹ and generative adversarial networks (GAN)²⁰ along with optimization techniques like Bayesian optimization and reinforcement learning.

VAEs are capable of learning a continuous space representation of molecules which can then be optimized to get molecules with target properties through techniques like Bayesian optimization and swarm optimization. These techniques are architecture agnostic and can be applied with different forms of VAE like junction tree VAE by Jin et al., grammar VAE by Kusner et al. and more^{21–24}. A VAE model was also paired with reinforcement learning for generating molecules with high binding affinity to a given target by Boitreaud et al.²⁵ GANs are generative models that learn the probability distribution of the training data, and sampling from the distribution can then be used to generate synthetic data points. This model has also been applied to the generation of molecules with desirable properties in works by Cao et al., Prykhodko et al., Guimaraes et al. and Maziarka et al.^{26–29}

However, the common theme across all the enhanced molecular generation models is an optimization algorithm requiring an oracle to calculate the property the model is being optimized for. Some properties are easy to calculate, while others, like binding affinity, take significantly longer. Molecular docking is a non-convex optimization problem and can take up to ≈ 10 min for large molecules on a CPU. An alternative widely used is a machine learning predictor model that takes the molecule and target as an input and predicts the binding affinity^{30,31}. However, this also comes with the caveat that the predictor model heavily depends on the initial training data, and hence, for molecule generation pipelines, as the model dynamically moves in the chemical space, the type of molecules being sampled also changes dynamically. This leads to a phenomenon called model degradation in which the performance of machine learning models declines as time passes³². Though faster, using machine learning models to predict binding affinity can become highly inaccurate as molecules start being sampled from regions of the chemical space unseen in the initial training set. Hence, there is a requirement for a predictor model that can: (1) perform the closest to that of physics-based docking software or a computation-ally demanding free energy calculation, (2) model which can work with a small dataset by learning posterior distribution accurately, and, (3) to develop a framework (over the machine learning model) that learns how to predict the binding affinity as the generator navigates through the chemical space.

Active learning is a popular technique in machine learning for training predictor models on datasets that are expensive to label. Bayesian active learning using uncertainty sampling was introduced in Computer Vision for Object detection³³. More scalable and dynamic Active Learning approaches were introduced to improve training and network accuracy³⁴. Active learning in cheminformatics was employed to conduct high throughput virtual screenings in existing databases^{35–37}. Warmuth et al. uses support vector machines to mine data from an extensive collection of databases with ligands docked to a protein³⁸. Raschka and Kaufman et al. summarize AI-based research for GPCR bioactive ligand discovery focusing on Active Learning³⁹. Fujiwara et al.⁴⁰ employ active learning using query by bagging to find structurally diverse hits in large databases. Gentile et al. conduct scalable AI-based virtual screening with deep docking⁴¹.

In this study, we propose a generative model agnostic active learning framework that can be used to accurately predict binding affinities throughout the optimization process of any generator model. The framework uses a Gaussian process regression model, updated at regular intervals using new training data obtained as the generator model is optimized to generate molecules with high binding affinity. This architecture was validated by integrating it with the MoleGuLAR pipeline proposed by Goel et al. It was found that using this reduced the training time by 70% while maintaining high accuracy⁴².

Methods

This section describes the various components of the proposed dynamic predictor model, which can be used to replace the previously used slow docking tool and an inaccurate ML-based predictor model with an accurate ML-based predictor that learns new distributions with minimal data sampling. Figure 1 showcases the proposed predictor model using uncertainty-based active learning. Subsection *Gaussian process regressor* describes the formulation of Gaussian process regressor (GPR), the base predictor model. Initially, *k* points are sampled randomly from a database of drug-like molecules, their binding affinities with the target are calculated, and the predictor is trained to predict the binding affinity of these molecules with the required target. The GPR is updated dynamically using active learning, detailed in the subsection *Active learning*. The GPR also returns the uncertainty in its prediction, and the molecules used for retraining the model are picked using an uncertainty threshold which is also dynamic. Choosing this threshold is detailed in the subsection *Dynamic uncertainty threshold*.

Gaussian process regressor

Gaussian process regressor (GPR) are the predictors used for molecular property prediction⁴³. GPR calculates the probability distribution for fit over all possible functions that fit the data for a given distribution. When the



Figure 1. Architecture of the proposed dynamic predictor. An incoming molecule is used as input to a Gaussian Process Regression model which returns a prediction and uncertainty (standard deviation, STD). If the uncertainty is above a given threshold, the ground truth value of the property is calculated and added to a repository. If k points are accumulated in the repository, the model is re-trained, and the uncertainty threshold is updated.

GPR predicts the chemical property of a new incoming molecule, the prediction is tractable, and a normal distribution is obtained with mean and covariance. Hence, not only can a GPR return a prediction, but it can also return the uncertainty associated with that prediction in the form of standard deviation. In GPR, we first assume a Gaussian process prior f(x) using a mean function m(x) and covariance kernel function k(x, x') represented by

$$f(x) \sim GP(m(x), k(x, x')) \tag{1}$$

The prior assumes a multivariate distribution depending on the input dimensions. The mean function is usually a constant or the mean of the input dataset. The covariance function can be of any form of a function as long as it satisfies the properties of a kernel. The kernel function used for this study is the radial basis function (RBF) kernel⁴⁴ represented by

$$k(x, x') = \sigma_f^2 exp(\frac{-1}{2l^2} ||x - x'||^2)$$
⁽²⁾

with hyper-parameters: signal variance (σ^2) and lengthscale (l). The GPR must be pre-trained with an initial set of points before inducting it into the Molecule generation pipeline. The input to the GPR for training is represented using SMILES. These smiles are converted into Mol2Vec embeddings⁴⁵. These Mol2Vec embeddings are fed into a pre-trained GPR to predict the desired molecular property.

Active learning

Active learning is a particular case of machine learning. Active learning algorithms can interactively query a database, acquire new data points, append them to the existing dataset, and retrain the current model. Active learning is used when unlabelled data is abundant, but labeling them is expensive and time-consuming - this is ideal. The function with which samples are acquired is known as the acquisition function. Common acquisition functions include balanced exploration, variance reduction, etc.

In this study, we employ Active learning using an uncertainty sampling-based querying strategy to acquire molecules whose properties are not predicted confidently by the ML model. When a set of molecules are to be evaluated by the oracle, the GPR initially predicts its uncertainty in predictions on each of these molecules. Depending on the Standard Deviation threshold, a molecule is deemed to be "certain" or "uncertain" with respect to the GPR's prediction. Suppose the GPR is certain about the predictions of a molecule. In that case, the GPR makes the prediction, and the predictions are sent forward for the generative pipeline to calculate rewards. If the GPR is not certain about the prediction software, and the results are then sent to the generative pipeline to calculate rewards. When *k* points are accumulated in the repository, these points are concatenated with the pre-trained points, and the GPR is retrained. This is done to learn the newly explored range by the generative pipeline for the associated molecular property. This step is performed repeatedly, and the number of retraining steps depends on the number of uncertain points encountered by the GPR.

Dynamic uncertainty threshold

After every retraining step in the active learning pipeline, the model's mean uncertainty for the regions it was initially trained on fluctuates. Hence, using a constant standard deviation threshold over multiple re-training steps can lead to inaccurate classification of a prediction as "certain" or "uncertain" concerning the ML model. Suppose the mean uncertainty of predictions of the pre-trained data during a retraining step is greater than the





standard deviation threshold. In that case, the majority of the points are deemed uncertain and vice versa. Hence, it is only ideal for the standard deviation threshold to vary as the model's mean uncertainty on the pre-trained data varies during every re-training step. To vary the threshold, a test set is maintained. After every retraining step, the uncertainties of the GPR on the new test set are recorded, and a histogram of the uncertainties is plotted after dividing the uncertainties into k bins. The uncertainties of the first bin are noted. Given the established benchmark (uncertainties belonging to the first bin), when trained on adequate data points for a given data range, the GPR can predict with lower uncertainty for data points belonging to any other successive bins. Hence, the new uncertainty threshold is the mean of the second bin.

Dataset

The dataset used for this problem is the HTS collection by Enamine⁴⁶. The dataset comprises approximately 2 million ligands, which were then docked to the Tau Tubulin Kinase 1 (TTBTK1), an important target for neurodegenerative diseases like Alzheimer's⁴⁷. Binding affinities in this dataset range from – 12 to 0 kcal/mol. The ground truth has been generated by docking the molecules using Autodock-GPU⁴⁸. We follow the same docking methodology as Goel et al., and the molecules have been docked by following the same procedure for the 4BTK protein present in the S6: Docking Methodology, Supplementary Information for MoleGuLAR: Molecule Generation using Reinforcement Learning with Alternating Rewards⁴². For this problem, k points are randomly sampled from this dataset. Figure 2 depicts a histogram showing the distribution of the binding affinities of the molecules present in the dataset. As one can observe, the data gets more scarce in the higher binding affinity regions (regions < -8 kcal/mol). The scarcity of data in more negative binding affinity regions during virtual screening can be attributed to various factors. One reason is the presence of structural constraints and synthetic challenges associated with molecules that exhibit extremely negative binding affinities. Additionally, when utilizing a general-purpose database containing millions of ligands screened against a specific protein, it is possible that the chemical space explored within this database is more biased towards other target sets, resulting in limited coverage of the negative binding affinity regions. Furthermore, virtual screening alone cannot adequately explore newer chemical regions without the validation and input of a chemist who can suggest potential modifications for a set of promising molecules identified through virtual screening. These modifications, however, are often constrained by intellectual property considerations, which restrict the accessibility of data pertaining to higher binding affinity regions that hold greater potential for therapeutic applications. Hence, it is expected that a model trained on this data would have high error and uncertainty in the higher binding affinity regions.

Results and discussion

This section reviews the results obtained by testing different parts of the proposed oracle represented in Figure 3. Subsection *GPRs for predicting binding affinity* discusses the performance of a GPR model trained to predict binding affinities to the TTBK1 protein. The following subsection compares two techniques for selectively labeling more data during molecule generation. Subsection *Active learning integrated with MoleGuLAR* explores how using the proposed enhanced predictor model improves the efficiency and accuracy of performing docking calculations and a conventional ML-based predictor model, respectively.



Figure 3. Active learning integrated with the MoleGuLAR pipeline.

GPRs for predicting binding affinity

The GPR is trained on an initial data pool available from the Enamine dataset. Five thousand points are randomly sampled, and their Mol2Vec⁴⁵ embeddings are extracted and trained using the following kernel: an additive RBF kernel⁴⁴ with a length scale of 5.0 and White Noise kernel with default noise level (1.0). A 10,000-point test set is also sampled from the Enamine dataset to examine the accuracy of the dataset. The metrics achieved on the test set are: Mean Absolute error: of 0.452 kcal/mol, Mean squared error: of 0.378 kcal/mol, and R2 score of 0.87. Other models, including graph isomorphism networks, graph attention networks, and fully connected neural





networks on Word2Vec embeddings, were tested. While fully connected neural networks performed the worst with an MAE of 1.2 kcal/mol, MSE of 2.56 kcal/mol, and R2 of 0.68, and GINs performed well with an MAE of 0.64 kcal/mol, MSE of 0.92 kcal/mol, and R2 of 0.74, Graph based models and other deep learning models required 70–80k points to achieve this accuracy. GPRs achieved benchmark accuracy with as low as 5000 data points. Our main goal in the optimization process is to explore newer chemical spaces more accurately and do so in less time, and labeling more points means employing the Physics-based property prediction software more—indicating an additional cost. Moreover, GPRs also provide us a better and more trivial estimation of uncertainty due to their probabilistic framework and incorporation of priors. Meanwhile, in the case of graph models and deep learning models, we used Monte Carlo Dropouts as approximations for deep Gaussian processes for estimating uncertainty σ , by extracting the variance across predictions. With sparse data inputs and lesser data points, along with shorter message passing in the case of small molecules meant that uncertainties fluctuated during every run. It was only with GPRs that we obtained a common trend where uncertainties were high in lesser explored regions of the dataset (<– 8 kcal/mol), and were low in data-abundant regions (<– 3 kcal/mol and >– 7 kcal/mol). Figure 4 represents the ground truth versus the predicted graph for the test set. This pool is the initially labeled dataset for the active learning problem.

Active learning versus random sampling

As shown in the previous section, GPRs work well for predicting binding affinities to the given target, and the prediction uncertainty returned by the GPR can be leveraged to perform active learning. However, it is important to acknowledge the contribution of uncertainty sampling during chemical space exploration. To do so, we compare our uncertainty-based querying strategy with random sampling. The initial pool of training data consists of 500 data points from the enamine dataset. At every iteration, a GPR model is trained on the training data, and the 500 points from the entire Enamine dataset for which uncertainty is maximum are appended to the training set. Conversely, in the case of random sampling, the 500 points inducted into the training set are chosen at random. Figure 5 showcases the mean absolute error on a hold-out test set and shows that active learning outperforms random sampling and provides a better-performing predictor model.

Active learning integrated with MoleGuLAR

Active learning versus random sampling

The enhanced predictor model was integrated into the MoleGuLAR pipeline as the oracle. In MoleGuLAR, 500 molecules are generated to perform policy gradients and 100 more for evaluation during every iteration leading to 250 molecules in each iteration. The GPR makes a prediction for each of these molecules; for any molecule with uncertainty higher than the threshold, a docking calculation is performed, and the labeled molecule is added to a repository of molecules. As soon as the repository reaches a size of 300, these are inducted into the training set, the GPR is retrained, and the repository is reinitialized. Simultaneously, another repository is maintained for comparison in which randomly chosen molecules are inducted. Figure 6 compares the MAEs of the two GPRs with training data sampled using different strategies. The figure shows that the binding affinity increases as the number of points increases, which is the converse of what was seen in the previous section. This can be attributed to the generator model moving in previously unseen regions of the chemical space for which representation is low in the training data. It can also be noticed that at ≈ 6750 points, the difference gets significant. The sudden increase in mean absolute error (MAE) can be attributed to the RL framework's initial exploration



Figure 5. MAE versus Number of points in the training set. 500 Molecules from the Enamine dataset with the most uncertain predictions are appended to the training set and the GPR is retrained and the MAE is calculated on the hold-out test set.





of new regions. During this exploration phase, the model encounters data from these regions, combined with a lack of uncertainty sampling, which hinders the accurate learning of distributions in the newly discovered space. As the gradients within the RL framework gradually decrease, the rate of exploration of new chemical spaces also slows down. Still, we see that the error in the case of Active Learning with GPR increases less steeply than random sampling and does not see large fluctuations.

Binding affinities of generated molecules

To analyze and compare the "quality" of generated molecules based on the choice of Oracle, 500 molecules were generated using the generator before optimization. Following this, the generator was optimized for more negative binding affinity using a pre-trained static GPR as the oracle and the proposed dynamic predictor model as

the GPR. At the end of the optimization process, 500 molecules are generated from both, and their ground truth binding affinities are calculated using AutoDock-GPU. The distribution of these binding affinities is shown in Figure 7. It is visible that using the proposed predictor leads to the generation of molecules with high binding affinities compared to using a static predictor. The reason for this is that static predictor has extremely poor quality predictions for molecules with high binding affinities and hence, fails to predict those values. Therefore, the proposed dynamic predictor model leads to better performance than a static predictor.

Course correction

Analysis was also performed to check how the quality of predictions changes when the GPR model is retrained and whether it learns more information about the region of the chemical space being sampled. To do this, 1000 molecules were generated using MoleGuLAR after the optimization process, and their binding affinities were calculated using AutoDock-GPU. Predictions are then made using two models - the model used in the pipeline before retraining and after retraining. The correlation between these predictions and the ground truth values are present in Figures 8a,b, respectively. It is visible that the predicted binding affinities in regions where binding







Figure 8. Course correction graph (a) before and (b) after re-training.

Oracle	Time
Autodock-GPU	43 h
Active learning	13 h
Pre-trained	8 h

Table 1. Time analysis: MoleGuLAR.

.....

Oracle	Number of points labeled inside / outside the pipeline
Autodock-GPU	25000 / 0
Active Learning	2600 / 5000
Pre- trained	0 / 5000

Table 2. Labeling Analysis: MoleGuLaR.

affinity < -10 kcal/mol are not close to the ground truth binding affinities. At the same time, after retraining, they lie much closer to the y = x line. There is also a significant improvement in the *R*2 score and the MAE.

Improvement in efficiency

With evidence from the previous sections that the dynamic predictor model performs better than a static MLbased predictor model, the next step is to see if the accuracy trade-off by using a predictor model leads to a significantly shorter run time. Shorter run times not only promote cost-effectiveness but also provides researchers to run multiple models simultaneously under different parameters. We use Autodock GPU in all our docking calculations since Autodock GPU provides us with an accelerated framework for docking calculations. With Autodock-Vina, on average, one molecule would take 20 s to dock. With Autodock GPU, we are able to accelerate this process to 3–4 s per molecule. We compare the time taken by three different versions of MoleGuLAR under three different settings: the base model proposed by Goel et al., the pre-trained model proposed by Goel et al., and our active learning framework. The time taken to optimize MoleGuLAR for 100 iterations and the total number of docked molecules is presented in Tables 1 and 2, respectively.

It is evident that in terms of time taken the pretrained model is the most efficient but it makes bad predictions. On the other hand AutoDock-GPU makes accurate predictions but can take up to 2 days and hence, is extremely time consuming. The pre-trained model labels 20% of total molecules compared to the number of molecules generated by the pipeline. The active learning also labels 20% outside the pipeline—just like the pre-trained model, but also labels an additional 10.4% inside the pipeline. Hence, the active learning-based predictor finds a balance between the quality of predictions and the time taken.

Conclusion

In this study, a solution is presented to make the de novo generation of drug-like molecules more efficient. Active learning and uncertainty sampling are used to reduce the execution time of molecule generation pipelines. The approach is validated by conducting rigorous experiments which test the accuracy and the correctness of the Active learning pipeline and showed that Active learning acts as a trade-off between complete docking and a pre-trained Machine learning model which explores a local non-linear function to learn about the binding pocket. We also show that the trade-off proves to be very important in improving the accuracy and the distribution of the predicted molecules in the pre-trained model. Further work can include reducing the number of labeled points to a greater extent and altering graph-based machine learning models to work with smaller datasets. But, using a simplistic base model in this problem significantly improves the execution time and reduces the number of docking calculations.

Data availability

The data, code, analysis, models and the generated molecules have been included at https://github.com/deval ab/Enhanced-MoleGuLAR.

Received: 7 March 2023; Accepted: 16 September 2023 Published online: 29 November 2023

References

- 1. Irwin, J. J. et al. Zinc20-a free ultralarge-scale chemical database for ligand discovery. J. Chem. Inf. Model. 60, 6065–6073 (2020).
- 2. Mendez, D. et al. Chembl: Towards direct deposition of bioassay data. Nucleic Acids Res. 47, D930-D940 (2019).
- 3. Fox, S., Farr-Jones, S. & Yund, M. A. High throughput screening for drug discovery: Continually transitioning into new technology. J. Biomol. Screen. 4, 183–186 (1999).
- 4. Zhu, T. et al. Hit identification and optimization in virtual screening: Practical recommendations based on a critical literature analysis. J. Med. Chem. 56, 6560–6572 (2013).
- 5. Maia, E. H. B., Assis, L. C., de Oliveira, T. A., da Silva, A. M. & Taranto, A. G. Structure-based virtual screening: From classical to artificial intelligence. *Front. Chem.* 8, 343 (2020).

- 6. Mehta, S. et al. MEMES: Machine learning framework for enhanced MolEcular screening. Chem. Sci. 12, 11710–11721 (2021).
- Gentile, F. et al. Deep docking: A deep learning platform for augmentation of structure based drug discovery. ACS Cent. Sci. 6, 939–949. https://doi.org/10.1021/acscentsci.0c00229 (2020).
 - 8. Glavatskikh, M., Leguy, J., Hunault, G., Cauchy, T. & Da Mota, B. Dataset's chemical diversity limits the generalizability of machine learning predictions. J. Cheminform. 11, 69 (2019).
- 9. Reymond, J.-L. The chemical space project. Acc. Chem. Res. 48, 722-730 (2015).
- 10. Lyu, J. et al. Ultra-large library docking for discovering new chemotypes. Nature 566, 224-229 (2019).
- 11. Devi, R. V., Sathya, S. S. & Coumar, M. S. Evolutionary algorithms for de novo drug design—A survey. Appl. Soft Comput. 27, 543–552 (2015).
- 12. Kerstjens, A. & De Winter, H. LEADD: Lamarckian evolutionary algorithm for de novo drug design. J. Cheminform. 14, 3 (2022).
- Zadorozhny, K. & Nuzhna, L. Deep denerative models for drug design and response. ArXivhttps://doi.org/10.48550/arXiv.2109. 06469 (2021).
- Bagal, V., Aggarwal, R., Vinod, P. K. & Priyakumar, U. D. MolGPT: Molecular generation using a transformer-decoder model. J. Chem. Inf. Model. 62, 2064–2076 (2022).
- 15. Gupta, A. et al. Generative recurrent networks for DE novo drug design. Mol. Inform. 37, 1700111 (2018).
- 16. Grisoni, F., Moret, M., Lingwood, R. & Schneider, G. Bidirectional molecule generation with recurrent neural networks. J. Chem. Inf. Model. 60, 1175–1183 (2020).
- 17. Bongini, P., Bianchini, M. & Scarselli, F. Molecular generative graph neural networks for drug discovery. *Neurocomputing* **450**, 242–252 (2021).
- 18. Mercado, R. et al. Graph networks for molecular design. Mach. Learn. Sci. Technol. 2, 025023 (2021).
- 19. Kingma, D. P. & Welling, M. Auto-encoding variational bayes. In: Bengio, Y. & LeCun, Y. (eds.) 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings. (2014).
- 20. Goodfellow, I. et al. Generative adversarial nets. In Advances in Neural Information Processing Systems Vol. 27 (eds Ghahramani, Z. et al.) (Curran Associates Inc, 2014).
- Jin, W., Barzilay, R. & Jaakkola, T. Junction tree variational autoencoder for molecular graph generation. In: Dy, J. & Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning, vol. 80 of Proceedings of Machine Learning Research, 2323–2332. (PMLR, 2018).
- Kusner, M. J., Paige, B. & Hernández-Lobato, J. M. Grammar variational autoencoder. In: International Conference on Machine Learning, 1945–1954. (PMLR, 2017).
- Griffiths, R.-R. & Hernández-Lobato, J. M. Constrained bayesian optimization for automatic chemical design using variational autoencoders. Chem. Sci. 11, 577–586 (2020).
- 24. Winter, R. et al. Efficient multi-objective molecular optimization in a continuous latent space. chemsc 10, 8016-8024 (2019).
- Jacques, B., Mallet, V., Oliver, C. & Waldispuhl, J. Optimol: Optimization of binding affinities in chemical space for drug discovery. J. Chem. Inf. Model. 60, 5658–5666 (2020).
- De Cao, N. & Kipf, T. MolGAN: An implicit generative model for small molecular graphs. ICML 2018 workshop on Theoretical Foundations and Applications of Deep Generative Models. (2018).
- Prykhodko, O. et al. A de novo molecular generation method using latent vector based generative adversarial network. J. Cheminform. 11, 1–13 (2019).
- Guimaraes, G. L., Sánchez-Lengeling, B., Farias, P. L. C. & Aspuru-Guzik, A. Objective-reinforced generative adversarial networks (organ) for sequence generation models. https://doi.org/10.48550/arXiv.1705.10843 (2017).
- 29. Maziarka, Ł et al. Mol-cyclegan: A generative model for molecular optimization. J. Cheminform. 12, 1-18 (2020).
- Öztürk, H., Özgür, A. & Ozkirimli, E. Deepdta: Deep drug-target binding affinity prediction. *Bioinformatics* 34, i821-i829 (2018).
 Nguyen, T. *et al.* Graphdta: Predicting drug-target binding affinity with graph neural networks. *Bioinformatics* 37, 1140-1147 (2021).
- 32. Mauri, L. & Damiani, E. Estimating degradation of machine learning data assets. ACM J. Data Inf. Qual. 14, 1-15 (2022).
- 33. Gal, Y., Islam, R. & Ghahramani, Z. Deep bayesian active learning with image data. In *Proceedings of the 34th International Confer*ence on Machine Learning, ICML'17, Vol. 70 1183-1192. (JMLR.org, 2017).
- Haussmann, E. *et al.* Scalable active learning for object detection. 2020 IEEE Intelligent Vehicles Symposium (IV) 1430–1435. (2020).
 Graff, D. E., Shakhnovich, E. I. & Coley, C. W. Accelerating high-throughput virtual screening through molecular pool-based active learning. *Chem. Sci.* 12, 7866–7881 (2021).
- 36. Warmuth, M. K. *et al.* Active learning with support vector machines in the drug discovery process. J. Chem. Inf. Comput. Sci. 43, 667–673 (2003).
- Ding, X. et al. Active learning for drug design: A case study on the plasma exposure of orally administered drugs. J. Med. Chem. 64, 16838–16853 (2021).
- Warmuth, M. K., Rätsch, G., Mathieson, M., Liao, J. & Lemmen, C. Active learning in the drug discovery process. In: Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic, NIPS'01, 1449-1456. (MIT Press, 2001).
- Raschka, S. & Kaufman, B. Machine learning and AI-based approaches for bioactive ligand discovery and GPCR-ligand recognition. *Methods* 180, 89–110 (2020).
- 40. Fujiwara, Y. *et al.* Virtual screening system for finding structurally diverse hits by active learning. *J. Chem. Inf. Model.* **48**, 930–940 (2008).
- Gentile, F. *et al.* Artificial intelligence-enabled virtual screening of ultra-large chemical libraries with deep docking. *Nat. Protoc.* 17, 1–26 (2022).
- Goel, M., Raghunathan, S., Laghuvarapu, S. & Priyakumar, U. D. Molegular: Molecule generation using reinforcement learning with alternating rewards. J. Chem. Inf. Model. 61, 5815–5826 (2021).
- 43. Williams, C. & Rasmussen, C. Gaussian processes for regression. In: *Advances in neural information processing systems* 8, 514–520. Max-Planck-Gesellschaft (MIT Press, 1996).
- Chang, Y.-W., Hsieh, C.-J., Chang, K.-W., Ringgaard, M. & Lin, C.-J. Training and testing low-degree polynomial data mappings via linear svm. J. Mach. Learn. Res. 11, 1471–1490 (2010).
- Jaeger, S., Fulle, S. & Turk, S. Mol2vec: Unsupervised machine learning approach with chemical intuition. J. Chem. Inf. Model. 58, 27–35 (2018).
- 46. Enamine. HTS collection. https://enamine.net/compound-collections/screening-collection/hts-collection (n.d). Accessed 23 November 2021.
- Sato, S. et al. Spatial learning impairment, enhanced cdk5/p35 activity, and downregulation of nmda receptor expression in transgenic mice expressing tau-tubulin kinase 1. Soc. Neurosci. 28, 14511–14521 (2008).
- Santos-Martins, D. et al. Accelerating autodock4 with gpus and gradient-based local search. J. Chem. Theory Comput. 17, 1060–1073 (2021).

Acknowledgements

We would like to express our sincere gratitude to Kohli Center on Intelligent Systems and IHub-Data, IIIT-Hyderabad, and DST-SERB (CRG/2021/008036) for their support of this research. Their resources were instrumental in allowing us to conduct this study and achieve our goals.

Author contributions

The authors of this manuscript, K.V. and M.G., collaborated on writing the code and preparing the manuscript. In addition to his involvement in the code and manuscript preparation, K.V. also created all the figures. The problem statement was developed by M.G., S.L., and D.P. . G.V. also made a valuable contribution to the study through his thorough review and insightful input. All authors reviewed the manuscript to ensure its quality and accuracy, working together to effectively execute the study and accurately report its results.

Competing Interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to U.D.P.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

© The Author(s) 2023