



OPEN ACCESS

EDITED BY

Zhi-Bei Qu,
Fudan University, China

REVIEWED BY

Wenjun Xie,
University of Southern California,
United States
Qiang Shao,
Chinese Academy of Sciences (CAS),
China

*CORRESPONDENCE

Junde Li,
✉ jul1512@psu.edu

RECEIVED 07 February 2023

ACCEPTED 12 May 2023

PUBLISHED 01 June 2023

CITATION

Li J, Beaudoin C and Ghosh S (2023),
Energy-based generative models for
target-specific drug discovery.
Front. Mol. Med. 3:1160877.
doi: 10.3389/fmmed.2023.1160877

COPYRIGHT

© 2023 Li, Beaudoin and Ghosh. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Energy-based generative models for target-specific drug discovery

Junde Li*, Collin Beaudoin and Swaroop Ghosh

Department of Computer Science and Engineering, Pennsylvania State University, University Park, PA, United States

Drug targets are the main focus of drug discovery due to their key role in disease pathogenesis. Computational approaches are widely applied to drug development because of the increasing availability of biological molecular datasets. Popular generative approaches can create new drug molecules by learning the given molecule distributions. However, these approaches are mostly not for target-specific drug discovery. We developed an energy-based probabilistic model for computational target-specific drug discovery. Results show that our proposed TagMol can generate molecules with similar binding affinity scores as *real* molecules. GAT-based models showed faster and better learning relative to Graph Convolutional Network baseline models.

KEYWORDS

energy-based models, drug discovery, graph neural networks, generative models, target-specific

1 Introduction

Since the dawn of the genomics era in the 1990s, drug discovery has gone through a transition from a phenotypic approach to a target-based approach (Swinney and Anthony, 2011). Most drug targets encoded by human genomes are complex multimeric proteins whose activities could be modified by binding with drug molecules (Overington et al., 2006). A ligand compound is a substance that forms a complex with the binding site of a protein target, if they are structurally complementary, for therapeutic effects (see Figure 1). The navigation in the molecule space to find molecular compounds with high binding affinity is called target-specific *de novo* drug discovery.

Traditionally, the ligand was initially identified by screening libraries of commercially available compounds, which are sequentially docked against the protein target. This ligand discovery and optimization process could be time-consuming and resource-consuming with lower probabilities of success (Keserü and Makara, 2009). Computational approaches effectively accelerate nearly every stage of drug development. Most computational approaches are based on generative machine learning models, such as Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) (De Cao and Kipf, 2018; Li and Ghosh, 2022). However, these generative models hardly work for target-specific drug discovery since they merely learn the molecular distribution.

A few computational target-specific approaches also exist in the literature. For instance, Gupta et al. (2018) developed a generative RNN-LSTM model to produce valid SMILES strings and fine-tuned the model with drugs with known activities against particular protein targets. Unfortunately, such prior knowledge of protein binders is sometimes unavailable especially for newly identified targets. A recent work in (Grechishnikova, 2021) released this constraint by framing target-specific drug design as a machine translation problem. However, this non-generative model design only provides a probabilistic mapping from targets to ligands, thereby failing to sample ligand candidates for drug targets. CogMol




FIGURE 1

Illustration of the protein-ligand pair with PDB ID 4O0B from PDBbind Database. The red dashed square indicates the cartooned binding site and the docked ligand. 4O0B corresponds to the novel HSP90 selective inhibitor which shows potential utility in treating central nervous system disorders. The figure was prepared with PyMol 2.5.2 (DeLano, 2022).

(Chenthamarakshan et al., 2020) combined a Variational Autoencoder network and a protein-ligand binding affinity regressor for generating ligand molecules. However, the loosely coupled components in CogMol make the sampling less efficient and overall architecture not target-specific. We developed a novel algorithm, Target-specific Generation of Molecules (TagMol), to efficiently sample ligand candidates for given drug targets. To our knowledge, TagMol is the first computational approach developed for target-specific drug molecule discovery in an end-to-end learning fashion.

TagMol adopts a protein-ligand binding affinity regressor, which assigns high energies for ligands incompatible with targets and low energies for those compatible. Thus, our approach falls within the theoretical framework of energy-based models (LeCun et al., 2006). Figure 2 illustrates the energy-based latent-variable generative TagMol model which consists of a protein encoder, a

ligand generator, a critic network (or discriminator) and an energy network which guides to finding ligands with higher binding affinity. TagMol iteratively generates and evaluates molecular ligands with the generator and discriminator until convergence is reached. For better learning representation with molecular graphs, graph Neural Networks (GNN) including Graph Convolutional Network (GCN) (Kipf and Welling, 2016) and Graph Attention Network (GAT) (Veličković et al., 2018) are adopted and compared for exploit deeper level of message passing. Apart from the GNN-based critic network, the energy network takes GNN layers as well for extracting richer graph features. As the latent variable z (in Figure 2) varies in the multivariate Gaussian distribution, the *fake* ligand prediction varies over the ligand set compatible with the protein target. The TagMol learning is supervised using critic loss and relative binding energy values. The energy network ensures that generated (or *fake*) ligands are compatible with protein targets, and the critic network guarantees they are as realistic as real molecules.

The contributions of this paper are three-fold: 1) We proposed a novel end-to-end energy-based generative model, TagMol, for target-specific drug discovery; 2) the ligand generator architecture incorporates an extra latent variable z which entails the generation of ligands with high binding affinity to the input protein target; 3) we implemented graph neural networks with attention mechanism and multiple relations that result in faster and better learning.

We cover the basics on generative models and energy-based models in Section 2, describe the overall TagMol algorithm, major components and corresponding cost functions in Section 3, present the experimental setup, ablation study and results in Section 4, and draw our conclusions in Section 5.

2 Background

The matching between protein targets and ligands are not unique nor one-to-one. As reported in (Chen and Shoichet, 2009), ten drug fragments screened from the ZINC small-molecule database (Sterling and Irwin, 2015) well inhibited the CTX-M structure, which is a new enzyme family for extended spectrum beta-lactamases. To exploit the deterministic and

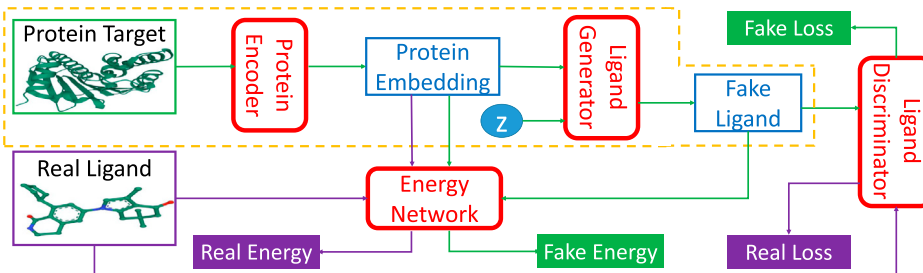


FIGURE 2

TagMol network architecture is composed of protein encoder, ligand generator, discriminator network, and the guiding energy network. Ligand generator, a latent-variable generative model, contains an extra latent variable z sampled from a multivariate Gaussian distribution. Energy network learns using energy differences between *real* and *fake* ligands. Blocks with green arrows indicate the generation flow of fake ligands; while blocks with purple arrows indicate real ligand workflow. After training, the network portion within the yellow dashed line can generate ligand candidates for a given protein target. The protein target and real ligand on the left are from the PDB 4O0B pair.

probabilistic model design benefits, we devised a latent variable energy-based model for drug discovery.

2.1 GAN-based models

Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) are implicitly generative models since they are evaluated using fake sample validity, predicted from a discriminator network. The generator of a GAN is a latent variable model with z being latent variables and x being observed variables. Conditional GAN (Mirza and Osindero, 2014) is an extended version of GAN which takes any auxiliary information, such as labels, into both the generator and discriminator. Based on conditional GAN, Barsoum et al. (2018) developed HP-GAN for probabilistic prediction of 3D human motions based on previous motions. Latent variables are necessary in modeling biomolecular PDBbind (Wang et al., 2004) refined 2017 dataset because the hidden target features, such as protein conformation and cellular localization, explicitly affect the formulation of small-molecule ligands. Based on the conditional GAN, TagMol takes as input the latent variables and protein targets for generating probabilistic ligand candidates for further screening. All possible atoms and bonds in the defined ligand space are assigned with certain probabilities in the generator accordingly. The latent variables would lead the predictions to different sets of plausible ligands conditioned on multiple protein families and conformations.

2.2 Energy-based models

Energy-based models (EBMs) (LeCun et al., 2006) capture dependencies between variables and evaluate their compatibility by associating a scalar *energy* value. The models are trained by designing an energy function which assigns low energies to correct pairs, and high energies to incorrect pairs. The loss function is designed to measure the quality of the energy function for assigning energy values to different variable pairs during learning and inference. The EBM framework covers a wide range of learning approaches, including probabilistic and deterministic, with respective loss functions. The discriminator in GAN is also an energy-based network which predicts the probability differences (energies) with zeros and ones for fake and real samples, respectively. The energy-based model for probabilistic prediction serves as the proxy for evaluating the binding energy between pairs of protein target and ligand. While the energy network is probabilistic, the protein encoder network is deterministic. As for the discriminator in GANs, the critic network in Figure 2 can also be considered as an energy-based network.

3 Approach

We explain in detail our probabilistic approach for target-specific drug discovery, conditioned on the given protein receptor in this section. The problem is defined as learning the conditional probability of plausible ligands $P(y|x)$, where $y \in \mathcal{Y}$, from a corresponding protein receptor $x \in \mathcal{X}$, given a training set of i.i.d.

protein-ligand pairs $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$. The ligand space \mathcal{Y} is composed of a bond adjacency matrix space $\mathcal{B} = \{0, 1\}^{N \times N \times B}$ and an atom matrix space $\mathcal{A} = \{0, 1\}^{N \times A}$, where N denotes the maximum number of heavy atoms (excluding Hydrogen) in ligand molecules; A and B represent the numbers of atom types and bond types, respectively.

3.1 TagMol algorithm

TagMol architecture is developed partially based on cGAN (see Algorithm 1). The generator creates synthetic (or *fake*) data samples from random noises, whereas the discriminator learns to distinguish between the *real* and *fake* samples. The adversarial minimax learning of cGAN is conditioned on extra information, such as class labels. Protein embedding serves as the conditional information in the present study. As depicted in Figure 2, the ligand prediction model takes as input a protein embedding x produced from the protein encoder, plus a latent vector z drawn from a Gaussian distribution. The protein embedding x and vector z are concatenated using early fusion and fed into a series of linear layers, as shown in Figure 3. The final atom layer and bond layer take the same fused features to generate probable atoms and bonds to form a possible ligand molecule. In our study, molecules are represented using graphs where each node denotes an atom and each edge denotes a bond. The following ligand discriminator (also called critic network since not trained to classify), represented with a Graph Neural Network (GNN) (see Figure 3), evaluates the generation quality. Generator and critic networks are the two major components in TagMol inherited from the GAN architecture. Apart from the evaluation from the critic network, the generated ligands should exhibit high hit rates when being docked with the provided protein target. To that end, a binding energy network (see Figure 2) is adopted to enforce target-specific generation.

Input: protein-ligand pairs p_{data} , iterations k , steps m
Parameter: network parameters τ_{Enc} , ϕ_G , ψ_D , θ_E , and hyper-parameters λ , α , β for loss terms
Output: predicted ligands \hat{y}

```

1: for  $k$  iterations do
2:   for  $m$  steps do
3:     Sample minibatch of protein-ligand pairs  $(x_p, y) \sim p_{data}$ .
4:     Get embedding from encoder  $x \leftarrow Enc_r(x_p)$ .
5:     Sample minibatch of noise samples  $z \sim p(z)$ .
6:     Generate fake ligands  $\hat{y} \leftarrow G_\phi(x, z)$ .
7:     Update D network parameters.
8:      $\psi_D := \arg \min \mathcal{L}_D(y, \hat{y}; \psi, \lambda) \triangleright$  see Eq. 2
9:   end for  $\psi$ 
10:  Repeat steps 3 to 6.
11:   $\mathcal{L}_E \leftarrow E_\theta(x, y) - E_\theta(x, \hat{y}) + \alpha \mathcal{L}_2 \triangleright$  see Eq. 9
12:  Update E network parameters.
13:   $\theta_E := \arg \min \mathcal{L}_E(x, y, \hat{y}; \theta)$ 
14:  Update  $\theta$  G and Enc network parameters.
15:   $\phi_G := \arg \min (-D(y) + \beta \mathcal{L}_E) \triangleright$  see Eq. 1
16:   $\tau_{Enc} := \arg \min (-D(y) + \beta \mathcal{L}_E)$ 
17: end for  $\tau$ 
18: return  $G_\phi(Enc_r(x_p), z)$ , for multiple  $z \sim \mathcal{N}(\theta, 1)$ 

```

Algorithm 1. Target-specific Generation of Molecules.




FIGURE 3

TagMol generator and discriminator components for ligand prediction. Protein embedding represents the extracted features from the input protein. A series of linear layers, atom layer and bond layer form the ligand generator. The ligand component indicates either a *real* ligand from PDBbind dataset or a *fake* one generated by TagMol. The *real* and *fake* ligands are to be evaluated by the following discriminator. A graph neural network forms the ligand discriminator, which assesses the prediction quality with the probability of generated ligand molecules being real.

3.2 Ligand generator and discriminator

TagMol is trained by first extracting a low-dimensional protein embedding space \mathbf{x} , as shown in Figure 2. The objective of the protein encoder $x = \text{Enc}_\tau(x_p)$ is to extract features associated with the protein binding pocket. A high-dimensional condition makes it hard for the model to build connections between generated ligands and complex proteins. An autoencoder-like unsupervised model learns the latent space representation for all protein targets, rather than the specific binding pockets of interest. Without adopting an autoencoder, the embedding network learns alongside all other components in an end-to-end fashion.

As indicated in Algorithm 1, G denotes the ligand generator and D the discriminator. Then the generated (or *fake*) ligand is represented as $\hat{y} = G_\phi(x, z)$ and discriminated with $D_\psi(y)$ where ϕ and ψ are learnable parameters in the generator and discriminator networks, respectively. The generator is a feed-forward neural network after fusing protein embedding and noise vector. While the discriminator is realized with either a GCN or GAT backbone for effectively learning graph representations. The baseline GCN is not specifically described since it is partially explained in GAT whose details are deferred to the separate Section 3.3. Each drawn latent vector z creates a plausible ligand molecule with different binding features for a protein target. Compared with VAE-based generative models, TagMol is susceptible to training instability (e.g., model collapse) as other general GAN-based models are. To mitigate such training issues, we replace the GAN with WGAN (Arjovsky et al., 2017) for measuring the approximation of generator distribution q to empirical distribution p with Earth Mover (EM) distance. Furthermore, a gradient penalty loss from WGAN-GP (Gulrajani et al., 2017) is adopted to enforce the WGAN Lipschitz constraint. The ligand generator is trained using WGAN-GP adversarial loss together with energy loss

$$L_G = -D(G(x, z)) + \beta \mathcal{L}_E. \quad (1)$$

The energy loss \mathcal{L}_E measures the docking energy difference between *real* and *fake* ligands. We remark that docking energy here is not computed based on atom interactions in terms of physical force fields, but on an energy function defined in Section 3.4. Different from cGAN, the discriminator only takes as input the *real* or *fake* ligands without concatenating the protein target condition. The relation between protein and ligand is guided by the energy network

which could be considered as another flexible discriminator. The discriminator loss consists of only the WGAN-GP critic loss

$$\mathcal{L}_D = D(G(x, z)) - D(y) + \lambda (\|\nabla_{\hat{x}} D(\hat{x})\|_2 - 1)^2 \quad (2)$$

where the interpolation $\hat{x} = \epsilon y + (1 - \epsilon)G(x, z)$ depends on a uniformly sampled weight $\epsilon \sim U[0, 1]$, and λ is a hyperparameter ($\lambda = 10$ is used in this study). The two terms to the left denote the WGAN loss and the right most term denotes the gradient penalty.

3.3 Graph neural network backbone

After a probabilistic distribution of molecules is produced from the generator, a hard categorical sampling step is realized using a straight-through trick for drawing a discretized one-hot ligand molecule represented by a bond matrix \mathcal{B} and an atom matrix \mathcal{A} . In this subsection, graph neural network backbone composed a series of GAT layers is explained to theoretically back up critic network and energy network.

As mentioned in Section 3.2, a set of GNN layers are adopted to learn graph-represented molecules by passing node messages iteratively. Bond types convey crucial information in formulating molecules and determine molecule valency validities. Therefore, the relational graph attention network (RGAT) (Qin et al., 2021) is specifically implemented for dynamically learning the importance of edge-specific attribute features. The input to a GAT layer is a molecule graph with $B = |\mathcal{R}|$ relation types and N nodes. The overall input node features are represented with a feature matrix $\mathbf{H} = [\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_N]^T \in \mathbb{R}^{N \times F'}$. The single-head attention coefficient for message passing is defined by incorporating multiple edge relations between i th and j th nodes

$$\alpha_{i,j}^{(r)} = \frac{\exp(\sigma(\mathbf{a}_r [\mathbf{W}_r \mathbf{h}_i \| \mathbf{W}_r \mathbf{h}_j]))}{\sum_{r' \in \mathcal{R}} \sum_{k \in \mathcal{N}_i^{(r')}} \exp(\sigma(\mathbf{a}_{r'} [\mathbf{W}_{r'} \mathbf{h}_i \| \mathbf{W}_{r'} \mathbf{h}_k]))} \quad (3)$$

$\forall i: \sum_{r \in \mathcal{R}} \sum_{j \in \mathcal{N}_i^{(r)}} \alpha_{i,j}^{(r)} = 1$, where $\|$ represents the concatenation operation, \mathbf{a}_r is the r -relation weight vector for the attention mechanism, $\mathcal{N}_i^{(r)}$ are all first-order neighbors of entity i with relation r , and σ is the LeakyRelu activation function used throughout our work. The attention weight $\alpha_{i,j}^{(r)}$ can be seen as the contribution from neighbor j to construct output node features \mathbf{h}'_i , after one GAT layer, represented as

$$\mathbf{h}'_i = \sigma \left(\sum_{r \in \mathcal{R}} \sum_{j \in n_i^{(r)}} \alpha_{i,j}^{(r)} \mathbf{W}_r \mathbf{h}_j \right) \quad (4)$$

where σ denotes the same LeakyReLU nonlinearity. The attention mechanism produces a single probability distribution over all neighbours of entity i irrespective of relation types. An output feature matrix $\mathbf{H}' = [\mathbf{h}'_1, \mathbf{h}'_2, \dots, \mathbf{h}'_N]^T \in \mathbb{R}^{N \times F'}$ is obtained with higher-order neighbor information. Multiple relational graph attention layers could be applied for learning better graph representation.

For ligand quality prediction, graph-level features are retrieved by referring to the graph aggregation method (Busbridge et al., 2019) which concatenates the mean of node representations with the feature-wise maximum across all nodes

$$\mathbf{g}(\mathbf{H}') = \left(\frac{1}{N} \sum_{i=1}^N \mathbf{h}'_i \right) \left\| \left[\begin{array}{c} F' \\ \oplus \\ \max_i \mathbf{h}'_{i,f} \end{array} \right] \right. \quad (5)$$

where \oplus denotes the element-level concatenation of feature maxima across nodes. A final validity multi-layer perceptron (MLP) neural network is concatenated for estimating the prediction quality in the discriminator network. Likewise, a final MLP network is added for predicting the binding affinities in the energy network.

3.4 Energy-based network

The probabilistic energy-based generative TagMol is developed by estimating the probability distribution $p(y|x)$ over the whole ligand space \mathcal{Y} for a certain protein x . Energy network aims to learn an energy function $E_\theta(x, y) \in \mathbb{R}$ that attributes low energies to regions near the data manifold $(x, y) \in \mathcal{X} \times \mathcal{Y}$ and high energies to other ligand regions. The energy function defines a probability distribution usually via a Gibbs-Boltzmann density

$$q_\theta(y|x) = \frac{\exp(-E_\theta(x, y))}{Z_\theta(x)}, \text{ and } Z_\theta(x) = \int \exp(-E_\theta(x, \tilde{y})) d\tilde{y} \quad (6)$$

where $Z_\theta(x)$ denotes the normalizing partition function. However, $Z_\theta(x)$ is generally intractable due to high dimensionality of target space \mathcal{Y} . Unlike Markov Chain Monte Carlo (Du and Mordatch, 2019; Nijkamp et al., 2020) to inefficiently approximate $Z_\theta(x)$, contrastive samples in this study are directly produced from the generator by referring to EBGAN (Zhao et al., 2017). The gradient of negative log-likelihood loss $\mathcal{L}_{\text{energy}}$ using contrastive samples is presented below:

$$\nabla_\theta \mathcal{L}_{\text{energy}}(\theta; p(x, y)) = -E_{p(x, y)} [\nabla_\theta \log q_\theta(y|x)] \quad (7)$$

$$\approx E_{p(x, y), z \sim p(z)} [\nabla_\theta E_\theta(x, y) - \nabla_\theta E_\theta(x, G_\theta(x, z))] \quad (8)$$

where $G_\theta(x, z)$ denotes the generated example from noise $z \sim p(z)$ with conditioning on protein embedding x_p . The trick from the last step is that the expectation w.r.t. \tilde{y} in Eq. 6 is approximated using a single contrastive example $\hat{y} = G_\theta(x, z)$ produced from the generator. The loss $\mathcal{L}_{\text{energy}} = E_\theta(x, y) - E_\theta(x, \hat{y})$ is still an object we want to minimize by pushing down the energies for *real* samples from the dataset and pulling up the energies for *fake* samples.

Equation 8 could also be interpreted as minimizing the density ratio between a pair of *fake* and *real* samples such that $Z_\theta(x)$ is bypassed. The final loss function for energy network is defined by adding a L2 regularization

$$L_E = L_{\text{energy}} + \alpha (E_\theta(x, y)^2 + E_\theta(x, \hat{y})^2). \quad (9)$$

4 Experiments and results

4.1 Dataset and metrics

All the experiments are conducted with the biomolecular PDBbind (Wang et al., 2004) refined 2017 dataset which contains 4,506 protein-ligand pairs. Ligands with more than 32 heavy atoms are trimmed by removing the atoms with a small number of bonds with the neighboring atoms. Heavy atom types include carbon, nitrogen, oxygen, fluorine, sulfur, and chlorine. Protein files with *.pdb* extension were loaded and pre-processed by removing hydrogen atoms. The 3D coordinates of protein atoms in angstroms plus atom types were extracted for representing the whole proteins instead of only binding sites. We remark that the number of atoms was set to 4,096 for all proteins to keep the same dimension for learning representation. Therefore, larger proteins were trimmed by removing atoms away from the corresponding ligand centroids, and smaller ones were padded with extra zeros.

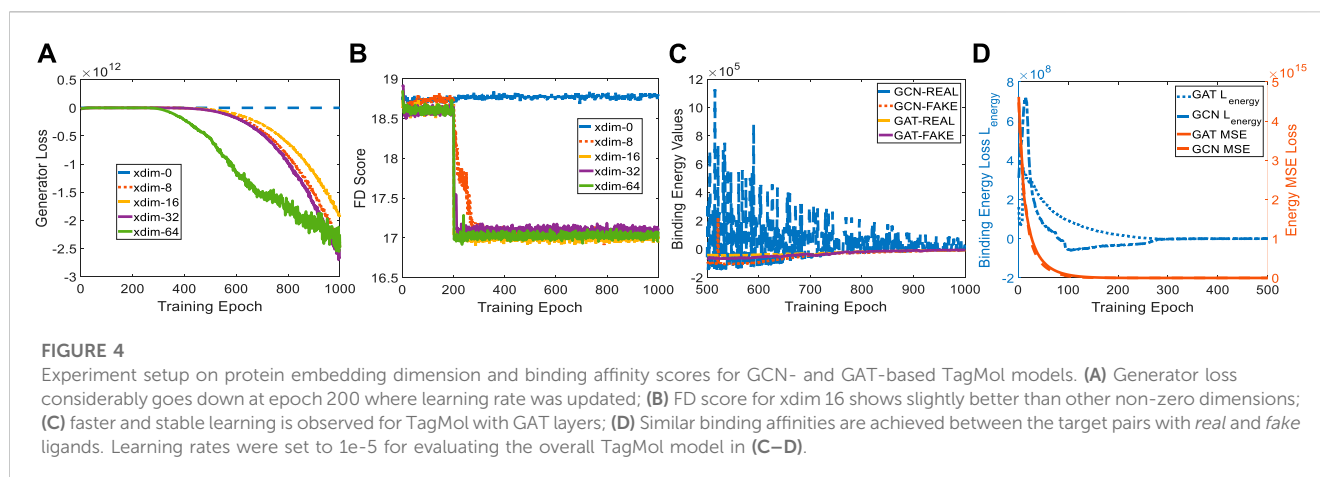
Learning results of the proposed cGAN-based models are evaluated using Fréchet distance (FD) which estimates the similarities between generated ligands and *real* ones. To evaluate the effectiveness of a protein target, variants without protein embedding ($x_{\text{dim}} = 0$) are trained as well for the purpose of comparison. Each sample batch of real or fake molecules is concatenated to a multidimensional point in the sampling distribution. Both atom and bond features from sampled molecules are considered for FD score calculation by referring to (Li et al., 2021). The performance from a non-conditional model without protein embedding serves as the FD baseline to evaluate the energy-based generative models.

4.2 Implementation details

Initially the energy network is dropped for conducting the ablation study on protein embedding dimension. All GAN variants are trained with a minibatch of 64 molecules with the Adam optimizer on a single RTX 2080 Ti GPU. The learning rate is initially set to 1e-4 and updated to 1e-5 after 200 training epochs. All models are trained with 1,000 epochs, and early stopping is applied if the learning diverges. Implementation details with the above dataset and metrics that support the findings of this study are available in the GitHub repository <https://github.com/jundeli/TagMol>.

4.3 Ablation study

The protein embedding dimension (x_{dim}) affects the cGAN performance in terms of generator loss. When x_{dim} is set to zero,



ligand generation is independent of the given protein, indicating a non-conditional model. Therefore, the binding pocket in the target cannot guide the target-specific drug discovery, as shown in Figure 4B where the FD score hardly decreases. When xdim is large, the model is more complex which corresponds to a steeper generator loss curve (see Figure 4A). However, the variance effect caused by Gaussian noise z is mitigated. The variance is beneficial since various ligands could be created for a certain target. Embedding dimensions ranging from 0 to 64 were tested for finding a suitable dimension that achieves a better FD score. A protein embedding dimension of 16 was selected through the ablation study on it. It is worth noting that the sudden changes in generator loss and FD score were caused by the learning rate decay at milestone epoch 200. All following experiments are conducted with 16 dimensional protein embedding.

4.4 Results

As mentioned earlier, no prior work has developed an end-to-end generative learning algorithm for target-specific drug discovery. Therefore, the result comparison is conducted by performing two TagMol variants with different GNN backbones in discriminator and energy networks. The energy-based network $E_{\theta}(x, y)$ reflects the final binding affinity between protein target and ligand candidates. The learning quality of generated ligands are thus evaluated using binding energy loss $\mathcal{L}_{\text{energy}}$ and scaled Mean Squared Error (MSE) loss $\alpha(E_{\theta}(x, y)^2 + E_{\theta}(x, \hat{y})^2)$. Hyperparameter α was set to 0.001 after several rounds of warm-up learning. Two types of GNN backbones, i.e., GCN and GAT, were tested for comparing the binding energy values. TagMol results with these two settings were plotted in Figure 4. We remark that, in panel (d), the final negative $\mathcal{L}_{\text{energy}}$ value reveals a better affinity for *real* ligands. The value eventually comes close to zero, which indicates the *fake* ligands have a similar binding affinity relative to *real* ones. The right y -axis displays the scaled MSE losses for GCN- and GAT-based energy models. A slightly lower MSE loss was observed for GCN models for the first few dozen epochs. $\mathcal{L}_{\text{energy}}$ and MSE curves become less distinguishable after learning 500 epochs.

To compare GCN and GAT in detail, we plotted the learning variance (or instability) of GCN-based energy models in Figure 4C.

The baseline GCN models showed relatively unstable curves due to the lack of an attention mechanism. The fluctuating loss for *real* ligands is possibly attributed to the bad GCN early-stage learning quality such that each weight update causes large binding energy changes. Binding energy values of GAT for *real* molecules turned out to be smaller than *fake* ones after 875 epochs. However, energy scores corresponding to GCN-REAL are still slightly higher than *fake* counterparts after 1,000 training epochs. This is another indicator of the advantage GAT layers have over GCN layers. We remark that the predicted energies provide the proxy for indirectly evaluating binding affinities, rather than physically evaluate the compatibility between protein-ligand pairs in terms of physical force fields, because of the log-likelihood loss function in EBMs. Likewise, the generated molecules are mostly disconnected due to the lack of a connection checking loss which penalizes the unconnected molecules. As shown in (De Cao and Kipf, 2018; Li and Ghosh, 2022), loss terms regarding the comparison between *fake* molecules and set-level *real* ones could not resolve the disconnectivity issue properly.

5 Conclusion

We proposed a probabilistic energy-based model called TagMol for target-specific drug discovery. The model specifically evaluates the binding affinity scores between protein-ligand pairs with an energy-based model. The protein embedding dimensions were tuned first within the default cGAN components for learning better generative representation. The energy network was then added and trained to enable target-specific drug discovery with binding affinity scores. Generated ligands achieved comparable binding energy scores for TagMol models with GAN and GAT layers. However, a faster and more stable learning is observed for GAT layers with attention over all atoms in drug molecules.

Data availability statement

The original contributions presented in the study are included in the article/Supplementary Material, further inquiries can be directed to the corresponding author.

Author contributions

JL developed the TagMol algorithm, trained and evaluated model performances, generated results and plots, wrote the manuscript. CB helped with Python implementation of TagMol and ran experiments. SG conceived the topic and critically read and edited the manuscript. All authors contributed to the article and approved the submitted version.

Funding

The work is supported in parts by NSF (CNS-1722557, CCF-1718474, OIA-2040667, DGE-1723687, DGE-1821766, and DGE-2113839) and seed grants from Penn State ICDS and Huck Institute of the Life Sciences.

References

- Arjovsky, M., Chintala, S., and Bottou, L. (2017). "Wasserstein generative adversarial networks," in International Conference on Machine Learning, International Convention Centre, Sydney, NSW, Australia, August 6–11, 2017 (PMLR) 70, 214–223.
- Barsoum, E., Kender, J., and Liu, Z. (2018). "Hp-gan: Probabilistic 3d human motion prediction via gan," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 1418–1427.
- Busbridge, D., Sherburn, D., Cavallo, P., and Hammerla, N. Y. (2019). *Relational graph attention networks*. *arXiv preprint arXiv:1904.05811*.
- Chen, Y., and Shoichet, B. K. (2009). Molecular docking and ligand specificity in fragment-based inhibitor discovery. *Nat. Chem. Biol.* 5, 358–364. doi:10.1038/nchembio.155
- Chenthamarakshan, V., Das, P., Hoffman, S., Strobelt, H., Padhi, I., Lim, K. W., et al. (2020). Cogmol: Target-specific and selective drug design for Covid-19 using deep generative models. *Adv. Neural Inf. Process. Syst.* 33, 4320–4332.
- De Cao, N., and Kipf, T. (2018). *Molgan: An implicit generative model for small molecular graphs*. *arXiv preprint arXiv:1805.11973*.
- DeLano, W. L. (2022). *The pymol molecular graphics system*. Available at: <https://pymol.org>.
- Du, Y., and Mordatch, I. (2019). Implicit generation and modeling with energy based models. *Adv. Neural Inf. Process. Syst.* 32.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2014). "Generative adversarial nets," in *Advances in neural information processing systems*.
- Grechishnikova, D. (2021). Transformer neural network for protein-specific de novo drug generation as a machine translation problem. *Sci. Rep.* 11 (1), 1–13. doi:10.1038/s41598-020-79682-4
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. (2017). Improved training of wasserstein gans. *Adv. neural Inf. Process. Syst.* 30.
- Gupta, A., Müller, A. T., Huisman, B. J., Fuchs, J. A., Schneider, P., and Schneider, G. (2018). Generative recurrent networks for de novo drug design. *Mol. Inf.* 37, 1700111. doi:10.1002/minf.201700111
- Keserü, G. M., and Makara, G. M. (2009). The influence of lead discovery strategies on the properties of drug candidates. *Nat. Rev. Drug Discov.* 8, 203–212. doi:10.1038/nrd2796
- Kipf, T. N., and Welling, M. (2016). *Semi-supervised classification with graph convolutional networks*. *arXiv preprint arXiv:1609.02907*.
- LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., and Huang, F. (2006). *A tutorial on energy-based learning*.
- Li, J., and Ghosh, S. (2022). "Scalable variational quantum circuits for autoencoder-based drug discovery," in *2022 design, automation & test in europe conference & exhibition (DATE)* (IEEE), 340–345.
- Li, J., Topaloglu, R. O., and Ghosh, S. (2021). Quantum generative models for small molecule drug discovery. *IEEE Trans. Quantum Eng.* 2, 1–8. doi:10.1109/tqe.2021.3104804
- Mirza, M., and Osindero, S. (2014). *Conditional generative adversarial nets*. *arXiv preprint arXiv:1411.1784*.
- Nijkamp, E., Hill, M., Han, T., Zhu, S.-C., and Wu, Y. N. (2020). On the anatomy of mcmc-based maximum likelihood learning of energy-based models. *Proc. AAAI Conf. Artif. Intell.* 34, 5272–5280. doi:10.1609/aaai.v34i04.5973
- Overington, J. P., Al-Lazikani, B., and Hopkins, A. L. (2006). How many drug targets are there? *Nat. Rev. Drug Discov.* 5, 993–996. doi:10.1038/nrd2199
- Qin, X., Sheikh, N., Reinwald, B., and Wu, L. (2021). Relation-aware graph attention model with adaptive self-adversarial training. *Proc. AAAI Conf. Artif. Intell.* 35, 9368–9376. doi:10.1609/aaai.v35i11.17129
- Sterling, T., and Irwin, J. J. (2015). Zinc 15–ligand discovery for everyone. *J. Chem. Inf. Model.* 55, 2324–2337. doi:10.1021/acs.jcim.5b00559
- Swinney, D. C., and Anthony, J. (2011). How were new medicines discovered? *Nat. Rev. Drug Discov.* 10, 507–519. doi:10.1038/nrd3480
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. (2018). "Graph attention networks," in *International conference on learning representations*.
- Wang, R., Fang, X., Lu, Y., and Wang, S. (2004). The pdbname database: Collection of binding affinities for protein-ligand complexes with known three-dimensional structures. *J. Med. Chem.* 47, 2977–2980. doi:10.1021/jm030580l
- Zhao, J., Mathieu, M., and LeCun, Y. (2017). "Energy-based generative adversarial networks," in *5th international conference on learning representations*. ICLR 2017.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The authors SG declared that they were an editorial board member of *Frontiers*, at the time of submission. This had no impact on the peer review process and the final decision.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.