Designing Optimized Drug Candidates With Generative Adversarial Network

Maryam Abbasi (✉ maryam@dei.uc.pt)

Univ Coimbra, Centre for Informatics and Systems of the University of Coimbra, Department of Informatics Engineering, Coimbra

Beatriz P. Santos

Univ Coimbra, Centre for Informatics and Systems of the University of Coimbra, Department of Informatics Engineering, Coimbra

Tiago C. Pereira

IEETA, Department of Electronics, Telecommunications and Informatics, University of Aveiro, Aveiro

Raul Sofia

Univ Coimbra, Centre for Informatics and Systems of the University of Coimbra, Department of Informatics Engineering, Coimbra

Nelson R. C. Monteiro

Univ Coimbra, Centre for Informatics and Systems of the University of Coimbra, Department of Informatics Engineering, Coimbra

Carlos J. V. Simões

BSIM Therapeutics, Instituto Pedro Nunes, Coimbra

Rui Brito

BSIM Therapeutics, Instituto Pedro Nunes, Coimbra

Bernardete Ribeiro

Univ Coimbra, Centre for Informatics and Systems of the University of Coimbra, Department of Informatics Engineering, Coimbra

Jose L. Oliveira

IEETA, Department of Electronics, Telecommunications and Informatics, University of Aveiro, Aveiro

Joel P. Arrais

Univ Coimbra, Centre for Informatics and Systems of the University of Coimbra, Department of Informatics Engineering, Coimbra

Research Article

Keywords: Drug Design, SMILES, Generative Adversarial Network, GAN, RNN

Posted Date: March 3rd, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1409047/v1>

License: This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

RESEARCH

Designing optimized drug candidates with Generative Adversarial Network

Maryam Abbasi^{1*}, Beatriz P. Santos¹, Tiago C. Pereira³, Raul Sofia¹, Nelson R. C. Monteiro¹, Carlos J. V. Simões², Rui Brito³, José L. Oliveira³, Bernardete Ribeiro¹ and Joel P. Arrais¹

*Correspondence:
maryam@dei.uc.pt

¹Univ Coimbra, Centre for Informatics and Systems of the University of Coimbra, Department of Informatics Engineering, Coimbra, Portugal
Full list of author information is available at the end of the article

Abstract

Drug design is an important area of study for pharmaceutical businesses. However, low efficacy, off-target delivery, time consumption, and high cost are challenges and can create barriers that impact this process. Deep Learning models are emerging as a promising solution to perform *de novo* drug design, i.e., to generate drug-like molecules tailored to specific needs. However, stereochemistry was not explicitly considered in the generated molecules, which is inevitable in targeted-oriented molecules. This paper proposes a framework based on Feedback Generative Adversarial Network (GAN) that includes optimization strategy by incorporating Encoder-Decoder, GAN, and Predictor deep models interconnected with a feedback loop. The Encoder-Decoder converts the string notations of molecules into latent space vectors, effectively creating a new type of molecular representation. At the same time, the GAN can learn and replicate the training data distribution and, therefore, generate new compounds. The feedback loop is designed to incorporate and evaluate the generated molecules according to the desired property at every epoch of training to ensure a steady shift of the generated distribution towards the space of the targeted property. The results demonstrate that the proposed framework can generate realistic, novel molecules that span the chemical space. The proposed Encoder-Decoder model correctly reconstructs 99% of the datasets, including stereochemical information. The model's ability to find uncharted regions of the chemical space was successfully shown by optimizing the unbiased GAN to generate molecules with a high binding affinity to the Kappa Opioid and Adenosine A_{2a} receptor. Furthermore, the generated compounds exhibit high internal and external diversity levels 0.88 and 0.94, respectively, and uniqueness.

Keywords: Drug Design; SMILES; Generative Adversarial Network; GAN; RNN

Introduction

In drug development, for a new drug to reach the final step and get approved, an estimated \$2.8 billion has been spent, and between 10 and 15 years of research were necessary [1, 2]. This is due to the fact that most drug candidates fail before reaching the last step of the process, with recent estimates pointing to a success rate of only 2% [3]. Such a low success rate implies that this is not just an expensive process but a high-risk one from a financial point of view, as most investments will fail. The high dimensionality of the chemical space has been identified as one of the main challenges, since it has been estimated that between 10^{33} and 10^{60} could be synthetically accessible [4, 5] and only a small fraction of this chemical space has been explored [6]. As the evaluation of the chemical space is a prohibitively expen-

sive process, it is crucial to find new strategies that can effectively narrow down the search space. Deep Learning (DL) methodologies have been gaining momentum as a promising solution for *de novo* drug design, whose goal is to generate novel molecular compounds that exhibit specific properties, such as being active towards a predefined biological target [7]. Two steps can be typically identified in this process: the first concerns creating a model that can replicate the chemical space. In contrast, the second focuses on optimizing the aforementioned model so that it is able to generate new molecules that exhibit specific properties.

Recurrent Neural Networks (RNNs) were the first DL-based technique that was successfully applied to drug generation. This type of network can learn and capture the syntax of sequences of data, which is the case of molecules, represented as Simplified Molecular Input Line Entry Strings (SMILES). Olivecrona *et al.* and Liu *et al.* resorted to this methodology to generate new molecules. They then optimized the model with the REINFORCE algorithm [8], a policy-based Reinforcement Learning (RL) method, in order to bias it towards the space of desired properties [9, 10]. Popova *et al.* approached this problem by pre-training two independent networks: a stack-augmented RNN Generator and a Predictor, which are then used jointly through RL to optimize the generator. They showed that their model could be effectively biased towards physical properties like the melting temperature and partition coefficient, specific biological activity, and chemical complexity [11]. Zheng *et al.* trained an RNN model with GRU cells on a biogenic dataset that includes stereo-chemical information to learn the grammar of these SMILES strings with higher complexity. They then fine-tuned it employing Transfer Learning (TL) [12]. Gupta *et al.* and Segler *et al.* drew similar conclusions by implementing the same method but using LSTM cells [13, 14]. The former also showed that its method could be applied to low-data drug discovery and fragment-growing while advocating that TL avoids introducing errors or unwanted bias compared to RL. The approaches mentioned above can suffer from exposure-bias [15, 16], which prompted the appearance of other DL-based alternatives for targeted generation of compounds, mainly adversarial approaches. A Variational Autoencoder (VAE) was first used in this context by Gomez-Bombarelli *et al.* to transform the SMILES discrete data into a real-valued continuous vector [17]. In this work, a model that can predict specific properties from the latent vector was also trained. Blaschke *et al.* studied and compared VAEs and adversarial autoencoders and were able to show that both methodologies result in preservation of the chemical similarity of the input molecules [18]. Among proposed approaches in this topic, the RL-based approaches are the most common, but have a tendency to focus on local minima and therefore return very similar and sometimes duplicate molecules. Applying a Generative Adversarial Network (GAN) that includes optimization strategy is an alternative approach which increase the diversity inside the generated compounds and resolve the problem of exploitation exploration of the reinforcement technique.

Additionally, stereochemistry is of critical importance to drug design and action, since the structure of a drug molecule is critical in defining how it interacts with different biological components [19]. However, it has often been neglected in *de novo* drug design and filtered from the training dataset in the state of the art, as it decreases the percentage of validly generated molecules.

In this work, the aforementioned issues are addressed by proposing a framework to generate candidate drug molecules. The framework includes an autoencoder with compound SMILES embedding, feedback GAN with gradient penalty for sequence generation, and an LSTM property predictor for feedback to GAN. The method was applied to designing inhibitors for the kappa opioid receptor (KOR) and Adenosine A_{2a} receptor (ADORA2A). The results showed that: (1) latent representations of SMILES were valid because over 99% of compounds (with or without stereochemical information) were reconstructed properly using the autoencoder; (2) The newfound representation is then used as real data to train the GAN to replicate the chemical space. GAN-generated compounds had similar distributions compared to original data, including predicted properties; (3) the Predictor can successfully evaluate the newly generated compounds by predicting their affinity for the desired target; (4) the feedbackGAN model can successfully be optimized by generating novel and diverse molecules including stereochemical information with high binding affinities towards the targets.

Methods

The general proposed framework is illustrated in Figure 1 and comprises an autoencoder, more specifically, an encoder-decoder architecture based on RNNs [20, 21], a Wasserstein GAN with gradient penalty (WGAN-GP) [22], a Predictor and an optimization step based on feedbackGAN [23]. The encoder-decoder architecture allows the model to learn a context vector (Figure 1-C) which is a new fixed-length representation of SMILES strings that can be mapped back to the original molecule. Therefore, the encoder (Figure 1-A) can be used to transform an entire dataset of

SMILES into an equivalent dataset that consists of context vectors. This new equivalent dataset is used as real data to train the WGAN-GP so that, once trained, the generator (Figure 1-D) is able to generate new samples that follow the same distribution as the context vector's dataset. In order to obtain the corresponding SMILES strings, these samples must then be passed through the decoder (Figure 1-B). By combining a GAN with an autoencoder, it becomes possible to train the WGAN-GP, surpassing the differentiation problem associated with discrete data, such as SMILES strings [24]. In this manner, the generator is able to generate molecules that span the entire chemical space. However, in drug discovery, the goal is to generate compounds that exhibit desired properties. In order to do so, an optimization step based on feedbackGAN [23] was devised that slowly forces the generator to focus on a specific region of the chemical space. Unlike the previous steps, this is a goal-specific step, as it requires a Predictor (Figure 1-F) for the desired characteristic that can identify interesting compounds. These compounds are then included in the training data by replacing the worst scoring compounds. This leads to a slow but effective shift in the generated distribution towards the desired region of the chemical space.

Encoder-Decoder Model

The Encoder-Decoder model, also known as a sequence to sequence model, is an autoencoder that contains recurrent layers, which allows it to work with sequences of data [25]. This type of model is typically used when the goal is to predict a new sequence of words that is in some way related to the input sequence, such as language translation tasks. The Encoder-Decoder learns a new representation of the input data, denoted by the context vector. This work uses this architecture to find a new continuous representation of SMILES strings that can then be used as training data for a GAN model, overcoming the inability to backpropagate through discrete data that arises if a GAN is trained directly on the SMILES String. The encoder translates the input SMILES string into a context vector belonging to the latent space while the decoder, given this context vector, reconstructs the initial SMILES Strings. Therefore, the model is trained using the categorical cross-entropy between the input SMILES and the predicted SMILES as the loss function. The optimizer uses the gradient of this loss function to update the network's weights.

The proposed Encoder-Decoder model is represented in Figure 2 resulted from the evaluation of several architectures and numerous hyperparameters. The Supplementary file contains a summary of these parameters and a detailed study of the efficiency of the model through grid search strategy. The optimized model consists of an encoder module with an embedding layer as its input layer, followed by two bidirectional LSTM layers with 512 units. The final hidden and cell states are retrieved from the bidirectional layers and concatenated before going through a dense layer (with as many units as the desired length of the context vector, 256 in this case) and a Gaussian Noise layer (standard deviation of 0.1). This last layer is only active during training and allows the context vector to become more robust. The decoder uses the context vector to feed four independent dense layers tasked with creating the hidden and cell states that will be given as initial states to its two stacked LSTM layers (with 512 units) followed by a dense layer. This dense layer includes a *softmax* activation layer in order to output the probabilities associated with the next token. Also, in between every one of the aforementioned layers, in both the encoder and decoder, a batch normalization layer is included (with batch normalization momentum of 0.9), which speeds up training and reduces overfitting by normalizing the output of the previous layer [26].

Wasserstein GAN with Gradient-Penalty

We adopted Wasserstein GAN with Gradient-penalty (WGAN-GP) since this type of model exhibits increased performance and stability [22] when compared to the traditional GAN as proposed by Goodfellow et al. [24]. GANs cannot be directly applied to categorical data, like SMILES strings, due to the fact that the sampling process at the end of the generator does not allow for the back propagation of the errors through that layer. However, since the proposed encoder-decoder model can find an alternative continuous representation for the SMILES strings, this new latent space representation can be used as real data to train the WGAN-GP without requiring intrinsic changes to the model. As the training data now consists of vectors, both the Critic and the Generator can be simple Feed Forward Neural Networks. These two neural networks have opposing objectives, with the critic aiming at distinguishing between generated $\tilde{x}$ and real data x , and the generator trying to fool the critic into believing that its data is real. A gradient penalty (weighted by λ) is included in the loss function (L) to prevent the critic's gradient from deviating from 1. Equation 1 shows this loss function, where the third term corresponds to the gradient penalty.

$$L = \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}} [D(G(\mathbf{z}))] - \mathbb{E}_{\mathbf{x} \sim p_{data}} [D(\mathbf{x})] + \lambda \mathbb{E}_{\hat{\mathbf{x}} \sim p_{\hat{\mathbf{x}}}} \left[(\|\nabla_{\hat{\mathbf{x}}} D(\hat{\mathbf{x}})\|_2 - 1)^2 \right]$$

If $\tilde{\mathbf{x}} = G(\mathbf{z})$ then $\hat{\mathbf{x}} = \epsilon * \tilde{\mathbf{x}} + (1 - \epsilon) * \mathbf{x}$ with $0 \leq \epsilon \leq 1$ (1)

where ϵ is uniformly sampled between 0 and 1. In practice, the gradient penalty amounts to creating a set of interpolated samples $\hat{\mathbf{x}}$ (each using a real sample $\mathbf{x}$, a generated sample $\tilde{\mathbf{x}}$ and ϵ) that result from randomly choosing points that lie on the lines that connect the batch of real samples to the batch of fake samples and evaluating its gradients. The gradient penalty loss term returns the squared

distance between the gradient calculated at the interpolated points and 1. Therefore, it penalizes the critic whenever its gradient deviates from 1, thus enforcing the 1-Lipschitz constraint [22].

The implemented WGAN-GP model includes a critic network and a generator network. The critic contains two dense layers with 256 units each and Leaky-Relu as the activation function (with $\alpha = 0.3$) and a final dense layer (256 units) with no activation function. The generator comprises five dense layers with Leaky-Relu activation function ($\alpha = 0.3$): the first contains 128 units and the remaining layers contain 256 units each; a batch normalization (momentum of 0.9) is included between the aforementioned layers. The input to the generator is a 64 dimensions vector drawn from the uniform distribution. Both networks are trained using the Adam optimizer with a learning rate of 0.0001 and drop-out value of 0.2.

Case-study

In this study, we used Kappa Opioid Receptor (KOR) and A_{2a} receptor (ADORA2A) as case studies. The Kappa Opioid Receptor is one of four opioid receptors that belong to the G-protein-coupled receptors (GPCR) superfamily. Research on opioid receptors, particularly on KOR, has been gaining momentum as it mediates affective disorders such as depression and anxiety, neurological diseases like epilepsy, but also pain and drug addiction, making it a promising pharmacological target [27, 28, 29, 30]. The ADORA2A is also a part of the GPCR superfamily and mediates pain, motor control, and mood [31]. The interest in ADORA2A, present in the central nervous, stems from the fact that it plays a crucial role in modulating motor functions, maintaining a low level of non-motor side effects. Therefore, ADORA2A antagonists are an attractive non-dopaminergic option for treating Parkinson's disease. The goal of the current experiment is to find ligands that bind to the KOR/ADORA2A and blocks their biological responses. To attain it, the generator of the WGAN-GP model must be optimized to generate compounds with a high affinity to bind antagonistically. The pIC_{50} is used to measure this, which is the negative logarithm of the half-maximal inhibitory concentration. Thus, the higher the pIC_{50} , the more potent the inhibition.

Predictor

In order to evaluate the optimization process, a QSAR model, henceforth denominated by Predictor, was implemented [32]. This regression model aims to predict the affinity of a given molecule for the KOR as measured by the pIC_{50} . The Predictor is an LSTM-based model that receives tokenized and padded SMILES strings as input, which are then passed through an embedding layer followed by two LSTM layers and two dense layers. Since this is a regression problem, the last layer has a single unit and a linear activation function. The use of RNN-based QSAR models is of particular interest for two reasons: first, this type of model works with inputs with different lengths, and second, it also works with SMILES strings, which means that there is no need to find other types of molecular representations that might add human bias [33]. The optimal architecture for the Predictor consists of an embedding layer with 128 units, followed by two LSTM layers and a dense layer (128 units each), and a final layer with a single unit (Figure 3).

Optimization through feedbackGAN

After training the implemented WGAN-GP model, the Generator is able to produce context vectors that are then decoded into SMILES strings that span the chemical space. In order to optimize the proposed framework, a feedbackGAN based approach is devised to bias the model towards specific properties. FeedbackGAN is an optimization framework proposed by Gupta et al. [23] that resorts to a feedback-loop and a function analyzer to optimize a GAN towards the space of desirable properties. In the context of the current problem, the Predictor takes the place of the function analyzer. After training the WGAN-GP model, the Predictor is linked to the GAN through a feedback mechanism. The GAN then enters a retraining phase in which, at the end of each epoch, the Generator is used to sample a set of new molecules fed to the Predictor to be evaluated. At each epoch, the n training SMILES with the worst score given by the Predictor are replaced by the n best scoring newly generated SMILES strings. By doing so, the training data is constantly being updated with new and better molecules according to the goal that has been set. This results in a fine-tuned Generator that gradually approaches the space of the desired property.

Validation Strategy

The validation strategy for the encoder-decoder model is based on the percentage of molecules that it can correctly reconstruct. As the training dataset only contains valid molecules, a correctly reconstructed (CR) SMILES string is automatically valid. The SMILES generated by the framework are syntactically and biochemically validated by RDkit tools [34], and also by four other metrics inspired by the Guacamole framework [35], which are the following: the percentage of valid molecules (Validity), the percentage of uniquely generated molecules (Uniqueness), i.e. non-repeating SMILES, the percentage of generated molecules that don't appear on a reference dataset (Novelty), and finally the Kullback-Leibler (KL) divergence between datasets. The KL divergence uses several properties of the molecules also calculated using the RDkit tools, namely: NumRotatableBonds, MolLogP, MolWt, BertzCT, TPSA, NumHAcceptors, MolWtNumHDonors, NumAliphaticRings, and NumAromaticRings. Besides these properties, it is also included the maximum nearest neighbor using ECFP4 fingerprints. Specifically, the KL divergence is calculated as presented by the following equation:

$$S = \frac{1}{k} \sum_i^k \exp(-D_{KL,i}), \quad (2)$$

where k is the number of calculated features, and $D_{KL,i}$ is the KL divergence calculated between the datasets of the i property's value distribution.

Finally, the *Internal Diversity* (Int Div) and *External Diversity* (Ext Div) are also computed. The diversity is evaluated by resorting to the Tanimoto Similarity, which computes the similarity between two molecules in terms of their circular fingerprints [36]. The Tanimoto distance can be defined as $1 - T_s$. From it, the diversity between two sets of generated molecules A and B can be defined as the average of the Tanimoto distance between every single pair of molecules:

$$Div(A, B) = \frac{1}{|A| \cdot |B|} \sum_{a \in A} \sum_{b \in B} (1 - T_s) \quad (3)$$

Therefore, the internal diversity is obtained by simply computing $Div(A, A)$ and the external diversity by $Div(A, B)$, where A is the set of generated SMILES and B is the set of the training data.

Results and Discussion

Datasets

To train the Encoder-Decoder and the WGAN-GP, a dataset containing SMILES strings from the ChEMBL [37] and Zinc Biogenic [12] databases is used, including, therefore, molecules with and without stereo-chemical information. For training and testing the predictor, besides the KOR and ADORA2A, we used the Ubiquitin specific protease 7 (USP7) and Tyrosine-protein kinase JAK2 (commonly called JAK2) and their corresponding $pIC50$ (see Table 1).

The SMILES strings are canonicalized, checked for duplicates, and then preprocessed. The preprocessing step consists of a character by character tokenization followed by adding 'G' at the beginning of each SMILES string and 'A' at the end and padding. The tokenized SMILES are then either One-hot Encoded (OHE), where each token becomes a binary vector or passed through an embedding layer that converts each token into a dense vector.

Once the best architecture and set of hyperparameters had been defined, the model was trained on two other more complex datasets: `composed_dataset.1` and `composed_dataset.2` which contain 100,000 and 500,000 drug-like molecules, respectively, that were retrieved from the remaining datasets mentioned in Table 1. This resulted in datasets that include a wider variety of compounds and molecules with and without stereochemistry.

Experimental Analysis on the Encoder-Decoder

The Encoder-Decoder model is trained using a batch size of 128 and the Adam optimizer (learning rate of 0.01). The maximum number of epochs is set to 100, but only the models with the best validation loss are considered (10 % of the data is used to validate the model). During the training phase, the teacher's forcing

Table 1 Summary of the datasets used throughout the experiment

Dataset	# Compounds	Labeled	Observations
ChEMBL [37]	1,178,946	No	
Zinc Biogenic [12]	108,283	No	
ADORA2A	4,729	Yes	CHEMBL251
KOR	5,262	Yes	CHEMBL237
JAK2	1,697	Yes	CHEMBL2971
USP7 [38]	1,109	YES	CHEMBL2157850
bbbp [39]	1,340	Yes	
composed_dataset_1	100,000	No	
composed_dataset_2	500,000	No

Figure 4 Comparison of the predicted pIC_{50} distributions for the original data and generated data by WGAN-GP model.

algorithm is applied so that the ground truth of the current step is given as input to the following step in order to minimize the propagation of errors. Table 2 shows the results obtained when training the aforementioned model with 500,000 SMILES Strings. It can be observed that the model effectively learned to map SMILES strings to the latent space and vice-versa, with a percentage of correctly reconstructed molecules of 99.2% and 99.0% for the training and test data set, respectively. It is important to note that correctly reconstructed (CR) molecules are automatically valid and that the validity is constantly higher, meaning that some molecules are reconstructed into valid molecules but not the same molecules that are given as input.

Table 2 Performance of the Encoder Decoder model based on percentage of Correctly Reconstructed (CR) molecules and validity of the train and test sets.

% CR (train)	% CR (test)	% Validity (train)	Validity (test)
99.2	99.0	99.9	99.8

Experimental Analysis on WGAN-GP

The implemented WGAN-GP model is trained on 100,000 molecules and for 10,000 epochs. The results of sampling 1,000 valid molecules using the whole framework are presented in Table 3 and Figure 4 from which it is possible to conclude that the GAN effectively learns the training data distribution as it generates data that follows that same distribution regarding the values of predicted pIC_{50} . It is worth noting that the validity of the generated data is only 30.2%. This low value is attributed to the fact that the WGAN-GP is learning to mimic a distribution of continuous vectors that are then converted into a series of discrete tokens that need to abide by certain rules to be considered valid SMILES strings, hence the difficulty in generating valid molecules.

Table 3 The performance of the WGAN GP model in comparison of the original data with the generated ones.

	Original Data	Generated Data
Maximum	9.974	8.454
Mean	6.003	5.984
Minimum	4.548	5.083
Standard Deviation	0.684	0.577
External Diversity	-	0.890
Internal Diversity	-	0.887
% Unique	-	100.0
% Valid	-	30.2

To further compare the generated data to the original training data in terms of general drug-like properties, Figure 5 and 6 show the relationship between the Quantitative Estimate of Drug-Likeness (QED) and Synthetic Accessibility (SA) score, and the relationship between the logarithm of the partition coefficient between n-octanol and water ($\log P$) and Molecular Weight (MW), respectively. By interpreting them, it is possible to conclude that the original and generated data are clearly overlapping, which means that the GAN successfully learned the training data distribution in terms of several of its properties. It should be noted that to make the plots more perceptive, instead of resorting to the 100,000 molecules in the training data, only 5,000 were used, but care was taken to make sure that the distribution of this smaller set in terms of predicted pIC_{50} was kept the same.

Performance Analysis of the Predictor

The Predictor module is trained to predict the binding affinity. During training, the Adam optimizer (with $\beta_1 = 0.9$ and $\beta_2 = 0.999$) with a learning rate of 0.0001 and a batch size of 16 is applied. The maximum number of epochs is set to 100, but early stopping is employed, and all the models stop training before reaching the 100th epoch. Figure 7 shows the scatter plot that results from applying the predictor to two different hold-out test sets, where it can be observed that the model performs constantly well for the complete range of values. Moreover, the Predictor module was assessed by acquiring the MSE, coefficient of determination, the concordance correlation coefficient (CCC), and Mathews correlation coefficient (MCC) on different datasets: USP7, JAK2, ADORA2A, and KOR. In the training phase, five

validation splits were used, and a model was created for each and independently trained. In the testing phase, each model tried to predict the test dataset and a mean of the results was calculated, as shown in (Table 4).

Table 4 Performance of the Predictor module on the USP7, JAK2, ADORA2A and KOR datasets.

Dataset	MSE $\pm$ Std	CD $\pm$ Std	CCC $\pm$ Std	MCC $\pm$ Std
USP7	0.065 $\pm$ 0.190	0.254 $\pm$ 0.126	0.431 $\pm$ 0.102	0.514 $\pm$ 0.048
JAK2	0.040 $\pm$ 0.003	0.259 $\pm$ 0.030	0.457 $\pm$ 0.036	0.253 $\pm$ 0.020
ADORA2A	0.021 $\pm$ 0.003	0.394 $\pm$ 0.026	0.587 $\pm$ 0.017	0.405 $\pm$ 0.042
KOR	0.026 $\pm$ 0.001	0.397 $\pm$ 0.012	0.583 $\pm$ 0.010	0.426 $\pm$ 0.030

Optimization with FeedbackGAN

The optimization has been done through feedbackGAN by maximizing KOR binding affinity. This process constantly updates the training data by replacing the worst scoring molecules with the best, which results in a gradual shift of the distribution towards the desired property. At the end of each epoch, 200 valid molecules are sampled, and 10% of the best replace the worst molecules in the “real data” training set. In order for the changes in the dataset to be significant and effectively bias the distribution that is being learned, only 5,000 molecules are used as “real data” for the optimization process. These 5,000 molecules are sampled from the dataset used to train the WGAN-GP, taking into consideration the distribution of the original dataset in terms of pIC_{50} . The unbiased WGAN-GP model is optimized for 500 epochs, considering that the best scoring molecules are the ones with the highest pIC_{50} . 1,000 valid molecules are sampled in intervals of 50 epochs and evaluated in terms of distribution of predicted pIC_{50} , validity, uniqueness, and diversity. These results are presented in Table 5. Figure 8 shows the distribution of predicted pIC_{50} values in terms of a probability density for intervals of 100 epochs of optimization, where it can be observed that the implemented strategy results in a successive shift of the overall distribution towards higher values of predicted pIC_{50} . The same conclusion can be drawn from Table 5 where it is clear that the mean of the distributions is constantly increasing from 5.984 for the unbiased model to 7.283 for the model optimized for 500 epochs. It is important to note that the

Table 5 Comparison of pIC_{50} distribution measures throughout the optimization process (maximization of KOR affinity). Max, Mean and Min value of pIC_{50} are shown in 500 epochs and for the unbiased (ub) model.

	ub	50	100	150	200	250	300	350	400	450	500
Max	8.454	8.341	8.852	9.340	9.354	8.911	9.149	9.000	8.939	9.241	9.179
Mean	5.984	6.501	6.719	6.791	6.851	7.015	7.074	7.168	7.204	7.273	7.383
Min	5.083	5.198	5.205	5.380	5.181	5.496	5.271	5.340	5.522	5.380	5.380
Int Div	0.887	0.890	0.890	0.887	0.891	0.886	0.885	0.885	0.883	0.878	0.877
Ext Div	0.890	0.900	0.909	0.908	0.915	0.923	0.922	0.926	0.930	0.937	0.938
% Uni.	100.0	100.0	100.0	100.0	100.0	100.0	99.9	100.0	100.0	99.8	100.0
% Valid	62.2	47.1	57.7	64.0	70.7	76.0	78.9	93.6	84.1	92.4	94.3

Figure 8 Distribution of the predicted pIC_{50} values for the unbiased model and the biased model (feedbackGAN) at every 100 epochs.

oscillations in the minimum and maximum values of the distributions are expected, as sampling is a stochastic process.

Although, as it has been shown, the proposed model was successful in maximizing the overall KOR affinity of the generated molecules, for them to be considered as potential drugs, it is also important to evaluate metrics such as QED, $\log P$, MW and SA score. Figure 9 and 10 show the scatter plots and distributions of these properties for 1,000 molecules generated after 500 epochs of optimization; the red square represents the region of drug-like properties. From Figure 10 it can be concluded that most of the molecules are outside the desired region (high QED and low SA score), with most of them exhibiting a SA score higher than 5, which implies that they are difficult to synthesize. It can be observed from Figure 9 that a significant amount of the generated molecules have good values for both $\log P$ and MW, even though there is a general tendency towards higher molecular weights. From Figure 10 it can be seen that as the pIC_{50} increases, there seems to be a tendency to generate more complex molecules, which is in agreement with the generally high SA score values. The supplementary file Section 5 provides some examples of the highest scoring molecules together with their properties for the sets generated with the unbiased model (at the top) and the optimized models at epochs 100, 200, 300, 400 and 500.

Experimental Analysis on Performance and Stereochemistry

We assess the performance of the model by analyzing the feedbackGAN mechanism's behavior given different molecule distributions as training. For the following experimental analysis, the ADORA2A molecules dataset was used. The pipeline generates the molecules based on two factors: the applied goals (i.e., for the selection of “best” molecules) and the input training dataset of the feedbackGAN.

From the ADORA2A dataset, three different ways of sampling 1000 molecules were performed, named: Random, Tani. Inf, and Tani. Sup [4]. Random sampling, as the name implies, selects molecules randomly. Tani. Inf and Tani. Sup selected the molecules with the highest and lowest internal diversity of the ADORA2A dataset, respectively. The internal diversity of each molecule was calculated using the Eq. 3. In order to avoid obtaining a final sampled pIC_{50} distribution too different from the original dataset, the molecules were sampled from arbitrary pIC_{50} intervals. Then the feedbackGAN was trained for 500 epochs for each sampled data and subsequently generated 10000 valid molecules. From Figure 11 it can be observed that, independently of the sampling method, the generated distributions of molecules all by a moderate extent converged to the same pIC_{50} values, but with different densities. In the case of the Tani. Inf, one potential hypothesis that could explain why the molecules are spread more sparsely is that the feedbackGAN might have difficulty generating molecules similar to sampled data (i.e., the 1000 sampled data), given that the molecules are very dissimilar between themselves. From this perspective, it makes sense that the sampling Tani. Sup generates molecules less sparse since the initial sampled data had lower internal diversity between molecules. Furthermore, the unbiased sampling of molecules (i.e., the random sampling) generates molecules with a distribution that falls between the other sampling methods (sparsity). It is important to highlight that pIC_{50} values of each molecule were calculated utilizing the trained predictor; therefore, the error might increase the further away the molecules are from the original dataset, which could also potentially explain the difference between distributions.

The molecules generated by each sampling approach are then evaluated by calculating the percentage of valid molecules (validity), the percentage of uniquely generated molecules (uniqueness), the percentage of generated molecules that do not exist

[4]The “Tani. Inf” and “Tani. Sup” correspond to the inferior and superior Tanimoto similarity, respectively, and not the Tanimoto distance.

Figure 11 Distribution of the predicted pIC_{50} values for different sampling methods from the ADORA2A dataset.

in the reference dataset (ADORA2A dataset) (novelty), and finally the Kullback-Leibler (KL) divergence between the generated molecules and the reference dataset, using several molecule descriptors, as implemented in the Guacamole benchmark [35]. The latter evaluation metric consists of calculating, for both datasets (i.e., the ADORA2A and generated datasets), several descriptors from the RDKit library, namely: NumRotatableBonds, MolLogP, MolWt, BertzCT, TPSA, NumHAcceptors, MolWtNumHDonors, NumAliphaticRings, and NumAromaticRings. On ECFP4 fingerprints, we also compute the distribution of maximal closest neighbor similarities. Each distribution of obtained values is compared between datasets using the KL divergence method, and a mean is calculated. The obtained values are presented in Table 6.

	Random	Tani. Inf	Tani. Sup
Validity (%)	82.81	89.89	93.22
Uniqueness (%)	96.89	94.73	97.63
Novelty (%)	100.0	100.0	100.0
KL Divergence	0.5266	0.2643	0.4915

Table 6 Performance analysis of generated molecules considering the sampling approach of the training dataset.

Note that both the Tani. Inf and Tani. Sup sampling methods generate (through the feedbackGan) a higher percentage of valid molecules than Random sampling. One possible explanation can be that inputting a very dissimilar (or similar) set of molecules can “force” the model to biased generate molecules in specific “drug spaces”, while the Random sampling method gives the model an unbiased exploration direction, which could explain why it has the lowest percentage of valid molecules. The KL divergence metric corroborates somewhat our premise that different sampling methods over the original dataset can compel the model to explore different “drug spaces”, as evidenced by the lowest value obtained by the Tani. Inf sampling comparatively to the other two methods.

Finally, a selection process was implemented to filter out molecules that do not have drug-like properties. Specifically, molecules that do not fall in the regions presented in Figures 10 and 9 are excluded. Additionally, a mean of the pIC_{50} values was calculated, and molecules with a value below the mean were excluded. This whole process was applied for the generated molecules of each sampling method. Figure 12 shows a sample of five molecules with the stereochemical information from each group. In chemistry, isomers can happen due to several reasons. In this

case, we are focused on stereoisomers, a form of isomerism that happens when two molecules have them arranged in a different, non-superimposable spatial orientation. These molecules have different geometries and therefore interact differently with the surrounding environment. A famous example of this phenomenon is Ibuprofen, where only one of the isomers of the active substance (sold as a mixture) is, in fact, active. Sometimes, however, two or more stereoisomers are very hard or even impossible to isolate or are interchangeable, and stereochemistry is generally omitted for these. But for carbon stereocenters, this is not the case, and that's the problem we address here. SMILES format has a form of depicting stereochemistry around each stereocenter. In double bonds, the characters `'/'` and `'\'` will denote orientation around it and only make sense if in pairs (for example, `/C = C/` will denote a cis and `\C = C/` a trans configuration). In chiral centers, the characters `'@'` or `'@@'` depict the local configuration of the ligands. However, they have no direct correspondence to R and S isomers until we consider the order of appearance of atoms in the SMILES string. This system, although being redundant, is not ambiguous.

The selected molecules from the three datasets contain all the stereochemical information they need to be fully defined in their SMILES. Note that all stereocenters have enough information not to leave doubts about what stereoisomer we are focusing on. This information is all encoded in each SMILES string, following the logic explained above. Predicted pIC50 is depicted in the legend. The greater this value, the higher the binding affinity. From all three datasets, and even from the sample shown here, it is possible to observe molecules sampled from Tani. *inf* (12) reveal the most structural variety, while the other two datasets seem to have adopted a "cycle plus chain" structure.

Technical notes: The implementations were in a computer with processors AMD Ryzen 9 3900X 12-core, 64 GB RAM and GPU Nvidia RTX 2070 8GB of GDDR6 VRAM using CUDA 10.1 with operating system Ubuntu 20.04.3 LTS. The results were coded in Python 3.8.3 using Tensorflow 2.3. The chemistry library used throughout is RDKit 2020.09.2 [34].

Conclusion

We propose a GAN-based framework that aims at generating and optimizing new molecules. As GANs cannot be straightforwardly applied to discrete data, which is the case of SMILES strings, we propose the use of an Encoder-Decoder model in order to obtain continuous representation of molecules. This model reached 98.8% of correctly reconstructed SMILES strings on a hold-out test set and 99.9% on the train set, which shows an improvement when compared to the state of the art. It is worth noting that our model is trained on a dataset that includes stereo-chemical information, which is often overlooked in most works due to its increased complexity but is critical in drug discovery. By passing the training data, that comprises SMILES strings, through the Encoder, an equivalent dataset is obtained that, instead of SMILES, contains fixed length continuous vectors. The representation is then used to train the GAN, more specifically a WGAN-GP. This type of model is able to implicitly learn the distribution of the training data and generate new

instances following that same distribution by pairing up two competing neural networks: a Generator and Discriminator. Once the WGAN-GP is trained, the Generator network is used to generate new vectors that are then converted into SMILES strings by the Decoder network. We showed that the trained WGAN-GP was able to replicate the distribution of the training data in terms of the predicted binding affinity for the KOR while still generating molecules with high levels of diversity (0.890 for the external diversity and 0.887 for the internal diversity). Even though the low percentage of valid generated molecules (30.2%) and the time required to train the models is clearly a drawback of this method, it is counterbalanced by the high diversity of the generated compounds and the percentage of their uniqueness (100%).

Once the WGAN-GP was trained on the dataset created by the Encoder, a feedbackGAN based optimization strategy is employed that consisted of continuing the training of the model and generating new molecules at every epoch. These new molecules are then evaluated according to their binding affinity for the KOR and the best scoring generated molecules replaced the worst scoring entries in the training data, resorting to a feedback loop. The obtained results proved that this strategy successfully resulted in shift of the generated distribution, with its mean moving from 5.984 for the unbiased model to 7.383 when aiming to maximize the predicted pIC_{50} of the generated molecules. Along with this, there was also an increase in the validity of the generated compounds from 30.2% to 62.3% with the internal diversity oscillating around 0.88 which implies that the model did not suffer from mode

collapse. Interestingly, the external diversity increased as the optimization process proceeded, meaning that the framework was indeed able to generate novel compounds with binding affinities as high as 9.18. In this sense, the devised framework was effectively maximized, though it should be noted that this was accompanied by an increase in the complexity of the generated compounds proved by the synthetic accessibility scores of the generated molecules.

The overall framework has the advantage that only the optimization step needs to be adapted to different problems and goals, simply requiring for that a problem-specific scoring metric function.

Availability of data and materials

All python code for this study is publicly available at <https://github.com/larngroup/GAN-Drug-Generator>.

Contact: maryam@dei.uc.pt

Competing interests

The authors declare that they have no competing interests.

Author's contributions

B.S. and M.A. implemented the model in the computational framework and analyzed the data; T.C.P. updated the model and created some complementary results and wrote some part of the manuscript; R. B. and C.S. help in understanding the different concept in biochemistry. R.S. helped analyze and select the stereochemistry information on the results and prepared figure 12 and wrote some part of the manuscript. M.A., J.P.A., J. L. and B. R. designed and directed the project; All authors discussed the results and contributed to the final manuscript.

Acknowledgements

This work is funded by national funds through the FCT - Foundation for Science and Technology, I.P., within the scope of the project CISUC - UID/CEC/00326/2020 and by European Social Fund, through the Regional Operational Program Centro 2020 and by the Portuguese Research Agency FCT, through D4 - Deep Drug Discovery and Deployment (CENTRO-01-0145-FEDER-029266).

Author details

¹Univ Coimbra, Centre for Informatics and Systems of the University of Coimbra, Department of Informatics Engineering, Coimbra, Portugal. ²BSIM Therapeutics, Instituto Pedro Nunes, Coimbra, Portugal. ³IEETA, Department of Electronics, Telecommunications and Informatics, University of Aveiro, Aveiro, Portugal.

References

1. DiMasi, J.A., Grabowski, H.G., Hansen, R.W.: Innovation in the pharmaceutical industry: new estimates of r&d costs. *Journal of health economics* **47**, 20–33 (2016)
2. Rashid, M.B.M.A.: Artificial intelligence effecting a paradigm shift in drug development. *SLAS TECHNOLOGY: Translating Life Sciences Innovation* **26**(1), 3–15 (2021)
3. Xue, H., Li, J., Xie, H., Wang, Y.: Review of drug repositioning approaches and resources. *International journal of biological sciences* **14**(10), 1232 (2018)
4. Polishchuk, P.G., Madzhidov, T.I., Varnek, A.: Estimation of the size of drug-like chemical space based on gdb-17 data. *Journal of computer-aided molecular design* **27**(8), 675–679 (2013)
5. Rodrigues, T., Reker, D., Schneider, P., Schneider, G.: Counting on natural products for drug design. *Nature chemistry* **8**(6), 531 (2016)
6. Mouchlis, V.D., Afantitis, A., Serra, A., Fratello, M., Papadiamantis, A.G., Aidinis, V., Lynch, I., Greco, D., Melagraki, G.: Advances in de novo drug design: from conventional to machine learning methods. *International journal of molecular sciences* **22**(4), 1676 (2021)
7. Schneider, G., Fechner, U.: Computer-based de novo design of drug-like molecules. *Nature Reviews Drug Discovery* **4**(8), 649–663 (2005)
8. Williams, R.J.: Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* **8**(3–4), 229–256 (1992)
9. Olivecrona, M., Blaschke, T., Engkvist, O., Chen, H.: Molecular de novo design through deep reinforcement learning. *Journal of Cheminformatics* (2017)
10. Liu, X., Ye, K., van Vlijmen, H.W.T., IJzerman, A.P., Van Westen, G.J.P.: An exploration strategy improves the diversity of de novo ligands using deep reinforcement learning: a case for the adenosine 2a receptor. *Journal of cheminformatics* **11**(1), 1–16 (2019)
11. Popova, M., Isayev, O., Tropsha, A.: Deep reinforcement learning for de novo drug design. *Science advances* **4**(7), 7885 (2018)
12. Zheng, S., Yan, X., Gu, Q., Yang, Y., Du, Y., Lu, Y., Xu, J.: Qbm-g: quasi-biogenic molecule generator with deep recurrent neural network. *Journal of cheminformatics* **11**(1), 1–12 (2019)
13. Gupta, A., Müller, A.T., Huisman, B.J.H., Fuchs, J.A., Schneider, P., Schneider, G.: Generative recurrent networks for de novo drug design. *Molecular informatics* **37**(1–2), 1700111 (2018)
14. Segler, M.H.S., Kogej, T., Tyrchan, C., Waller, M.P.: Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS central science* **4**(1), 120–131 (2018)

15. Guimaraes, G.L., Sanchez-Lengeling, B., Outeiral, C., Farias, P.L.C., Aspuru-Guzik, A.: Objective-reinforced generative adversarial networks (organ) for sequence generation models. *arXiv preprint arXiv:1705.10843* (2017)
16. Yu, L., Zhang, W., Wang, J., Yu, Y.: Seqgan: Sequence generative adversarial nets with policy gradient. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31 (2017)
17. Gómez-Bombarelli, R., Wei, J.N., Duvenaud, D., Hernández-Lobato, J.M., Sánchez-Lengeling, B., Sheberla, D., A., J., Hirzel, T.D., Adams, R.P., Aspuru-Guzik, A.: Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science* **4**(2), 268–276 (2018)
18. Blaschke, T., Olivecrona, M., Engkvist, O., Bajorath, J., Chen, H.: Application of generative autoencoder in de novo molecular design. *Molecular informatics* **37**(1–2), 1700123 (2018)
19. Renslo, A.: *The Organic Chemistry of Medicinal Agents*. McGraw Hill Professional, ??? (2015)
20. Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078* (2014)
21. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. *arXiv preprint arXiv:1409.3215* (2014)
22. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of wasserstein gans. *Advances in neural information processing systems* **30** (2017)
23. Gupta, A., Zou, J.: Feedback GAN (FBGAN) for DNA: a Novel Feedback-Loop Architecture for Optimizing Protein Functions (2018)
24. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. *arXiv preprint arXiv:1406.2661* (2014)
25. Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473* (2014)
26. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: *International Conference on Machine Learning*, pp. 448–456 (2015). PMLR
27. Schmidhammer, H., Erli, F., Guerrieri, E., Spetea, M.: Development of diphenethylamines as selective kappa opioid receptor ligands and their pharmacological activities. *Molecules* **25**(21), 5092 (2020)
28. Lalanne, L., Ayranci, G., Kieffer, B.L., Lutz, P.-E.: The kappa opioid receptor: from addiction to depression, and back. *Frontiers in psychiatry* **5**, 170 (2014)
29. Mores, K.L., Cummins, B.R., Cassell, R.J., Van Rijn, R.M.: A review of the therapeutic potential of recently developed g protein-biased kappa agonists. *Frontiers in pharmacology* **10**, 407 (2019)
30. Shang, Y., Filizola, M.: Opioid receptors: Structural and mechanistic insights into pharmacology and signaling. *European journal of pharmacology* **763**, 206–213 (2015)
31. Shang, Y., Filizola, M.: Opioid receptors: Structural and mechanistic insights into pharmacology and signaling. *European Journal of Pharmacology* **763**, 206–213 (2015)
32. Pereira, T., Abbasi, M., Ribeiro, B., Arrais, J.P.: End-to-end deep reinforcement learning for targeted drug generation. In: *In ICCBB 2020: 2020 4th International Conference on Computational Biology and Bioinformatics* (2020)
33. Chakravarti, S.K., Alla, S.R.M.: Descriptor free qsar modeling using deep learning with long short-term memory neural networks. *Frontiers in Artificial Intelligence* **2**, 17 (2019)
34. Landrum, G., et al.: RDKit: Open-source cheminformatics (2006)
35. Brown, N., Fiscato, M., Segler, M.H., Vaucher, A.C.: Guacamol: benchmarking models for de novo molecular design. *Journal of chemical information and modeling* **59**(3), 1096–1108 (2019)
36. Bajusz, D., Rácz, A., Héberger, K.: Why is tanimoto index an appropriate choice for fingerprint-based similarity calculations? *Journal of cheminformatics* **7**(1), 1–13 (2015)
37. Gaulton, A., Bellis, L.J., Bento, A.P., Chambers, J., Davies, M., Hersey, A., Light, Y., McGlinchey, S., Michalovich, D., Al-Lazikani, B., et al.: ChEMBL: a large scale bioactivity database for drug discovery. *Nucleic acids research* **40**(1), 1100–1107 (2012)
38. Wang, Z., Kang, W., You, Y., Pang, J., Ren, H., Suo, Z., Liu, H., Zheng, Y.: Usp7: Novel drug target in cancer therapy. *Frontiers in Pharmacology* **10** (2019)
39. Pereira, T., Abbasi, M., Oliveira, J.L., Ribeiro, B., Arrais, J.: Optimizing blood–brain barrier permeation through deep reinforcement learning for de novo drug design. *Bioinformatics* **37**(1), 84–92 (2021)

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [2022JournalBMCCHEminformaticssuplimentary.pdf](#)