

On the value of using 3D-shape and electrostatic similarities in deep generative methods

Giovanni Bolcato^a, Esther Heid^b and Jonas Boström^{*,c}

^aMolecular Modeling Section, University of Padova, Italy. ^bDepartment of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, United States ^cMedicinal Chemistry, Research and Early Development Cardiovascular, Renal and Metabolism (CVRM), BioPharmaceuticals R&D, AstraZeneca, Gothenburg, Sweden.

ABSTRACT

Multi-parameter optimization, the heart of drug design, is still an open challenge. Thus, improved methods for automated compounds design with multiple controlled properties are desired. Here, we present a significant extension to our previously described fragment-based reinforcement learning method (DeepFMPO) for the generation of novel molecules with optimal properties. As before, the generative process outputs optimized molecules similar to the input structures, now with the improved feature of replacing parts of these molecules with fragments of similar 3D-shape and electrostatics. By performing comparisons of 3D-fragments, we can simulate 3D properties while overcoming the notoriously difficult step of accurately describing bioactive conformations. We developed a new python package, ESP-Sim, for the comparison of electrostatic potential and molecular shape, allowing the calculation of state-of-the-art partial charges (e.g., RESP with B3LYP/6-31G**) obtained using the quantum chemistry program Psi4. The new improved generative (DeepFMPO v3D) method is demonstrated with a scaffold-hopping exercise identifying CDK2 bioisosteres. All code is open-source and freely available.

INTRODUCTION

A crucial task in all drug discovery projects is designing molecules against multiple, often contradictory objectives [1]. Much of today's drug hunters' time is therefore spent on attempting to find an optimal compromise where all desirable properties are satisfied in a single molecule. The use of sophisticated computational methods, leveraging high-quality datasets to help solve this task, is thus conceptually very attractive.

Recent advances in artificial intelligence (AI) and machine learning (ML) have given rise to an immense popularity of inverse design [2], and the field shows little signs of slowing down [3]. In inverse design, desired properties are specified *a priori*, and such methods generate compounds fitting that description [4]. Significant progress has been made in this area and a plethora of approaches for deep learning in molecular design has been published the last few years [2]. Many methods include reinforcement learning [5-7] to generate molecules, most often in the form of SMILES strings [8]. Other popular methods include generative methods such as recursive neural networks, generative adversarial networks or variational autoencoders, which are sometimes steered with reinforcement learning to control the molecular properties. The SMILES format in itself is nothing but amazing [9]. Using SMILES is also convenient for the AI algorithm, since a string is trivial to manipulate and transform. In addition, there are success stories of using SMILES in the area of generative design [10]. However, all molecules are

3D objects and a conservative modification to a SMILES string may cause a large effect in their 3D structure. Examples include the removal of brackets denoting substitution, such that a Y-shaped compound becomes linear, or the removal or changing of ring-closing locants. Therefore, optimization of molecular structures cannot be smooth in the space of 3D properties even though the SMILES strings change by only small amounts from iteration to iteration of the AI algorithm. We have previously presented a fragment-based generative approach (DeepFMPO) that addressed these modifications to the structure issue, albeit as 2D descriptions [11]. Here, we introduce a significant extension to DeepFMPO, using detailed descriptions of 3D-properties to represent molecules more accurately.

Shape and electrostatic properties of molecules are primary determinants of molecular recognition and should consequently be the method of choice when comparing the similarity of molecules encountered at various stages in drug design. Even though these methods have been used to achieve major impacts in related areas (e.g., virtual screening leading to the discovery of novel and unexpected chemotypes [12-14]), they have been largely unexplored in the context of *de novo* generative methods, although promising attempts have been made [15,16]. One reason for the reluctance of using 3D methods is the challenge of obtaining accurate descriptions of molecules bioactive conformations. In this work, we reduce this notoriously difficult step by using 3D-fragments of the complete target compounds.

Moreover, much of the work in the generative *de novo* design area has been focused on the development of maximally expressive methods whose purpose is to explore the entire chemical space. Our approach is different in this regard, since it specifically rewards the generation of molecules that are similar to known lead compounds. A related method is the recently published MolDQN method, which maximizes a “drug-likeness” (QED) score while also maintaining similarity to the original molecule [17]. To showcase the value of using shape and electrostatic similarities in deep generative methods, a case study highlighting its use in a scaffold-hopping exercise identifying bioisosteres for a set of CDK2 kinase inhibitors [18] is described.

METHODS

The DeepFMPO method is based on an actor-critic model for reinforcement learning [11]. It is a fragment-based generative method that learns how to modify compounds and improve them. That is, molecules are split into fragments, and these fragments are replaced with other similar fragments in the (deep) learning process of generating novel molecules with optimal properties. Technically, the fragments are encoded into binary strings, and similar fragments are assigned similar encodings. This is achieved by constructing a balanced binary tree. In the process of assembling the tree, similarities between all fragments are calculated. Fragments are paired in a greedy bottom-up manner, where the two most similar fragments are paired first. The joining procedure is repeated until all fragments are put together in a single tree. Subsequently this information is used to generate encodings for all fragments. The paths from the root to the leaves defines the encoding for each fragment. For every branch in the tree a one (“1”) is added to the encoding when going to the left and a zero (“0”) is added when going to the right, see Figure 1. The rightmost character in the encoding corresponds to the branching closest to the fragment. In this process, the pairwise similarity between all fragments is calculated. There are many ways to calculate chemical similarities, and the most used approaches currently employ 2D fingerprints.




Figure 1. A snippet of the balanced binary tree used in DeepFMPO. Fragments that are similar are placed close to each other. The encoding of a fragment is determined by the path from the root to the leaf. Every branching to the left adds a “1” to the end of the encoding and a branching to the right adds a “0”.

Here we present a new implementation of DeepFMPO utilizing a 3D-based molecular alignment method, where the ElectroStatic Potential (ESP) similarity between pairs of fragments is calculated. To this aim, we developed a python package, ESP-Sim, which calculates the overlap integrals of the electrostatic potentials (generated from Coulomb potentials) of two molecules or fragments. Within DeepFMPO, the computation of ESP similarities can be broken down into six steps for each fragment pair (see Figure 2a) and is described in more detail below. Steps 2, 3 and 6 correspond to function calls of the new ESP-Sim package, whereas steps 1 and 5 are innate to DeepFMPO. It is worth noting that this fragment alignment approach eliminates the challenging step of generating bioactive conformations for complete molecules, as well as alleviates the issue of aligning them correctly.

The Molecular Alignment of Fragments

All single bonds in a molecule that extend from a ring atom are broken in the DeepFMPO process, creating the molecular fragments. The attachment atoms (previously connected with a single bond) are labeled in this step. To calculate ESP similarities the fragments must be aligned in 3D. Here, a conformational search is conducted generating an ensemble of low-energy conformers for all fragments containing rotatable bonds, using the ETKDG method [19] as implemented in RDKit [20]. As default a maximum of ten conformations of each fragment is generated. An anchor group is connected to the fragments’ attachment atom and serves as a template in the alignment procedure. The coordinates of the anchor group are fixed in 3D space. The rationale for this step is that ligands containing related fragments typically bind in a similar orientation [21, 22], and these fragments will frequently make similar ligand-protein interactions. Consequently, to ensure that the fragments are aligned as accurately as possible, an anchor group is attached to the fragments and used in the molecular alignment step. The anchor group was arbitrarily chosen to be a hexazine ring with a methylene linker subunit. This group is of reasonable size for a template and highly unique (i.e., hexazines are never present in drug-like molecules) for easy identification and removal downstream in the process. A few experiments were conducted with other types of structural fragments as anchors to gauge possible conformational effects (*vide infra*). For each pair of fragments, the pair of conformations with the best shape overlay in terms of the highest shape Tanimoto value are stored. The anchor is then replaced with a hydrogen (see Figure 2). In cases where fragments include several labeled atoms, these are replaced with a methyl group. In this manner, all labelled atoms are replaced by a methyl which may be considered neutral in terms of electrostatic similarities. Finally, the ESP Tanimoto value is calculated between the pair of conformers with the best shape alignment (see the “*Electrostatic Similarity Calculations*” section below).




Figure 2. a) The stepwise procedure to obtain the electrostatic shape potential similarity values for pairs of fragments. b) An example of the corresponding procedure in graphics.

Electrostatic Similarity Calculations

The new ESP-Sim method uses the cheminformatics toolkit RDKit [20] to generate different conformations of a molecule with a constrained anchor or core group, and computes shape and electrostatic potential similarities between pairs of conformers. Alternatively, ESP similarities can be computed on pre-aligned molecules. The electrostatic potential similarity is computed via the overlap integral of the Coulomb potentials of two molecules, as well as their respective self-overlap integrals as

either Tanimoto [23,24] or Carbo similarity [25,26]. The Coulomb potential $V(\mathbf{r})$ describes the electric potential at a point \mathbf{r} as a sum of potentials of point charges q_i at points \mathbf{r}_i as

$$V(\mathbf{r}) = \frac{1}{4\pi\epsilon_0} \sum_i \frac{q_i}{|\mathbf{r} - \mathbf{r}_i|}$$

where ϵ_0 is the vacuum permittivity. Since analytic integration of the Coulomb potential at $\mathbf{r} = 0$ is not possible, we provide options to either approximate each potential with a sum of three Gaussian functions and integrate the fit function analytically analogous to Good *et al.* [27] or to perform a Monte-Carlo integration over the space outside of the van-der-Waals radii of each atom, and inside a user-defined margin. Partial charges can either be supplied by the user, calculated via RDKit (Gasteiger or MMFF94 charges), or computed using the open-source quantum chemistry program Psi4 [28], with the option of using the state-of-the-art Restrained Electrostatic Potential (RESP) charges [29]. There are a range of different methods and basis sets available in Psi4. For example, the often-recommended combination of using the B3LYP method and the 6-31G** basis set, although using those can be computational demanding. Within DeepFMPO, Dask [30], a library for parallel computing in Python, is used to speed up the process. It should be noted that the RESP/Psi4 method is not parametrized for atoms beyond the atomic number of Argon. To allow for larger atoms (e.g., bromine), their van der Waals (vdW) radius needs to be specified separately. In this code version, we set the vdW radii for bromo to 1.8 (file: resp/vdw_surface.py), following the GAMESS scheme [31] derived from the Merz-Kollman-Singh publication [32]. In addition to electrostatic similarities, ESP-Sim can furthermore output the shape Tanimoto similarity of molecules, describing the volume overlap. For DeepFMPO, we used the Tanimoto similarity of electrostatic potentials obtained via fitting to Gaussian functions (ESP-Tanimoto). We furthermore provide an option to add the volumetric shape score resulting in an ESP-TanimotoCombo score.

RESULTS

In the following retrospective case study, we aim to demonstrate the value of using shape and electrostatic similarities in scaffold-hopping exercises. Scaffold hopping is a method for identifying bioisosteric replacements [33,34] with the intention of retaining biological activity of analog compounds but also improving other relevant molecular properties. It can also be used as a design strategy for intellectual property (IP) reasons.

Assessing Various Molecular Similarity Measures

A frequently occurring scenario is that a drug hunting team has identified a promising compound, from an internal lead generation effort or from the literature, that needs optimization.

For the sake of argument, compound **1** in Figure 2 is such a compound [18].




Figure 3. Two equipotent CDK2 kinase inhibitors. CDK2 inhibitors containing the related bicyclic heterocycles imidazopyridine (**1**) and pyrazolopyridine (**2**) were discovered through high-throughput screening by Fischmann et al [18] and here used as a scaffold-hopping example.

With compound **1** at hand, the design question is then "which compound should we make next?". The optimization task usually includes improving molecular properties (e.g., permeability, solubility, clearance, selectivity, etc.) and perhaps also IP related issues. A common scenario then is for the project team, to try to come up with ideas of novel central rings to be introduced as scaffold replacements. In this context, it should be noted that heterocyclic rings are often considered special and typically end up in different patent applications [35]. Also, with regards to calculating molecular properties, many 2D-based methods are not adequately parametrized and have difficulties in assessing heterocyclic compounds accurately. So, how can break-through ideas for novel central rings be generated and which methods can be used to do it? Here, compound **2** (Figure 3) is one answer to the question "what to make next?". It is equipotent to compound **1** and importantly, contains a different but related central scaffold. That is, the bicyclic heterocycle in compound **1** (imidazo(1,2-a)pyridine) and compound **2** (pyrazolo(1,5-a)pyridine) are both 9-membered ring-systems with identical substituents.

To investigate how different methods predict the similarity of these kind of central bicyclic heterocyclic scaffolds we first generate a dataset of fragments containing the same framework and similar substitution pattern. Thus, the ChEMBL v28 database [36] was queried for compounds including a 9-membered bicyclic ring system, with three substituents, using SMARTS matching [20]. For comparison reasons, the substituents were subsequently removed providing 30 different scaffolds, see Figure 4. In this manner we identified an extensive list of 9-membered bicyclic heterocyclic scaffolds present in drug-like molecules that potentially could act as replacements for the pyrazolopyrimidine in compound **1**.

All 30 bicyclic systems were subsequently subjected to pair-wise comparisons using a range of standard 2D similarity measures, together with the new 3D-based similarity measure. A summary of the results obtained from each method is reported in Table 1. For completeness, the results using four different anchor fragments (hexazine, carboxylic acid, piperidine and iodine) are shown in Figure 5. The heat-maps are essentially the same

indicating that the method is not dependent on the choice of anchor fragment.

Table 1. Rankings for the **1** vs **2** fragment pair, among pair-wise comparison of 30 different heterocyclic rings. The rankings, and Tanimoto value, using a range of different 2D similarity methods available through RDKit and the new ESP-Sim measure are reported. Hexazine was used as anchor fragment.

Method	Rank (max = 30)	Tanimoto
ESP-Sim (B3LYP/6-31G**)	1	0.88
Morgan fingerprint (radius 2)	5	0.44
Morgan fingerprint (radius 3)	5	0.31
MACCS keys fingerprints	17	0.72
MCS Tanimoto	21	0.50
Topological fingerprints	22	0.23

The **1** vs **2** fragment pair is top-ranked when using the ESP-Sim (B3LYP/6-31G**) metric, but not by the 2D-based methods. The Morgan fingerprints rank the **1** vs **2** pair among the top five (Table 1), which is reasonably high. However, given the challenges and resource investments required to establish new synthetic routes, our experience is that very few alternative ring analogs are explored in real-life projects. Typically, only a couple of ring replacements are made and tested, essentially enforcing that only top-ranked scaffolds would be followed-up. Two other observations provide further support for the use of the ESP-Sim method. First, the MACCS keys fingerprint resulted in very similar values for many scaffolds (e.g., the Tanimoto similarity value for five scaffolds against scaffold of compound **1** show identical values - 0.87), suggesting that the MACCS keys similarity metric is not sufficient for capturing such subtle differences. Second, there are a couple of clearly structurally dissimilar fragments in Figure 4 (e.g. 1,4,6-trimethylpyrazolo[5,4-b]pyridine vs 2,4,7-trimethylimidazo[2,1-f][1,2,4]triazine) that are ranked low when using ESP-Sim (as they most probably should), but top-ranked when using Morgan 2D-fingerprints.




Figure 4. Bicyclic heterocyclic scaffolds in ChEMBL compounds matching the SMARTS pattern "[A][cH0]1[c,n][c,n]([A])[c,n]2[c,n][c,n][c,n]([A])[c,n]2[c,n]1".

As a final observation, deriving ESP-similarities with methods of lower theory for calculating the underlying partial

charges (Gasteiger, MMFF and HF/3-21G, data not shown) also yielded the 1 vs 2 pair as top-ranked, suggesting that such partial charges may be sufficient and a cost-effective alternative for this purpose.




Figure 5. All-against-all comparison experiments were conducted with four structurally different anchor fragments (top-left: hexazine, top-right: carboxylic acid, bottom-left: piperidine, bottom-right: iodine). The different anchors give essentially the same results.

Generating “Sweet Spot” Molecules

Having established the value of using the ESP-Sim measure (Table 1), the next step was to include it in the generative (DeepFMPO) method. An experiment was set-up to mimic a real-world scenario, where a set of lead compounds is optimized toward sweet spot criteria through a multi-parameter optimization process. Three different calculated properties (Molecular Weight, Polar Surface Area [20] and clogP [37]), commonly used in the optimization of leads to candidate drugs, were selected for this purpose. It should be noted that the choice of molecular properties was also selected for practical reasons facilitating reproducibility. Namely, there are methods to calculate them using RDKit [20]. The aim of the set-up was to bias the generation of compounds to fulfill the criteria for the three calculated properties, while also maintaining their similarity in shape and electrostatics towards a known set of lead compounds. The agent in the reinforcement learning method was rewarded for producing valid molecules and got a higher reward when generating molecules with properties in the targeted ranges. Since this was a scaffold-hopping exercise, with the goal of identifying a new bioisosteric scaffold, the minimum and maximum target values for the three properties were centered around the corresponding values for compound 1 (i.e., $320 < MW < 420$, $2.3 < \text{clogP} < 4.3$ and $45 < \text{PSA} < 65$).

The library of input fragments was generated from a set of structurally diverse compounds known to exhibit inhibitory effects against kinase targets, including compounds that have shown activity against the specific biological target of interest (CDK2). The data set was extracted from the ChEMBL database (version 28) using simple text searches, resulting in a set of 557

fragments (including the ones in Figure 4), as obtained from 1059 compounds. The lead series compounds were obtained by a substructure search using the (imidazo(1,2-a)pyridine) central scaffold of compound 1 on the surechembl website (<https://www.surechembl.org/search>) and yielded 138 close analogs, which is a typical number to what a drug hunting program would have access to. The data sets are available online (<https://github.com/giovanni-bolcato/deepFMPOv3D>).




Figure 6a-c. Graphs showing how molecular weight, logP and TPSA values change during the epochs, as the mean value of all the compounds for each epoch.

DeepFMPO with the ESP-Sim measure generated a total of 6359 unique molecules, when terminated at 1000 epochs. About two-thirds of those were sweet spot compounds. Hence, the agent generated compounds that have all three properties within the desired ranges. This number (ca 4000) is lower than when using a standard generative method facilitating the selection process, and a result of intentionally biasing using 3D-similarity. The evolution of the percentage of generated molecules that

demonstrate properties within the target ranges during the training process is shown in Figure 6a-c, displaying evidence of learning. A significant number of the generated compounds include the central scaffold of compound 2, and a number of those show near identical substitution pattern to compound 1. These bioisosteric compounds were observed in early epochs. One other nine-membered scaffold (3) was also represented among the generated, see Figure 7. When performing the same experiment but with simpler standard similarity measures (Morgan fingerprints, MACCS keys and Topological fingerprints), no compounds with the central scaffold of compound 2 (or 3) were generated. This provides an incentive for the use of DeepFMPO with ESP-Sim in scaffold-hopping exercises.

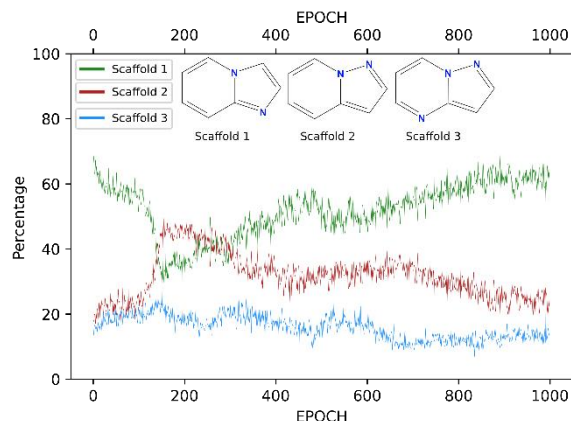


Figure 7. A graph showing the frequency of occurrence of compounds including the central fragment of compound 1, 2 and 3. The y-axis represents the total number of compounds each epoch (percentage).

DISCUSSION

In the current work we set out to explore the use of a new sophisticated similarity metric in generative methods. The power of rewarding compounds that are similar in 3D aspects, in addition to other molecular property constraints, is often underappreciated. It is a challenging task, due to the issues involved with conformer generation and molecular alignments. Nonetheless, this is a design strategy that we believe should be given more attention and we discuss why below.

Molecular representations in deep generative methods

Deep generative models typically use non-3D methods to represent molecules. Text-based methods and the use of SMILES strings is still the most prevalent representation. The reason for this is probably because SMILES can be massively expressive, and that it is trivial to manipulate and transform strings. However, there are some drawbacks with using SMILES strings [11, 38]. A significant problem is that a conservative change can have a huge change in the 3D structure of a molecule. This is important since all molecules are 3D objects. Here, we have addressed this issue by extending the fragment based DeepFMPO method, where molecules are built from similar fragments, instead of sequences of letters (as is the case for SMILES based methods). Fragment-

based methods are considered intuitive and often mimic the way medicinal chemists think and design. The approach was recently described by Meyers *et al* as a method “offer an appealing compromise between molecular expressivity and practicality” [38]. Hence, a common medicinal chemistry design strategy is to work on molecular series, swapping fragments and substituents in one part of the molecule, while keeping other parts of the compounds unaltered. This is often a challenge for generative methods working on SMILES strings [38], leading us to the next topic of discussion.

Deep generative methods can generate many compounds

Most generative AI methods produce tens of thousands of unique and diverse high-scoring compounds when used without stringent filters. This is related to Brenner’s underdetermined inverse problem stating that available data does not uniquely specify systems [39]. Also, although there may be nothing chemically wrong with AI generated molecules (i.e., all atoms in common valences and charge states), some can be exotic, and an experienced medicinal chemist would reject them upfront [40]. The issue of such unwanted molecules is manageable from a technical perspective. For example, one can enforce substructure rules and penalize the existence of undesired moieties (e.g., radicals, peroxides, anhydrides, strained and chemically unstable systems) in the reward functions, or as post-filters.

A more difficult problem to address is how to prune down the very many generated compounds to the few worth making. In reinforcement learning, a scoring function is used for this purpose. A complicating factor here is that drug discovery is complex and not all factors used in decision-making are easy to capture and thus not readily converted into rules that the AI methods can use in their rewards system. For example, a compound with several stereocenters is usually difficult to make (and resolve) and should consequently get a low reward score unless its building blocks are already available on the shelf. Also, absorption is a critical parameter for the optimization of oral drugs. Permeability over Caco-2 cells is often used as a surrogate when assessing absorption. A complicating factor here is that the uptake over the Caco-2 cells can be hampered by efflux, and in the case of high cell permeation the efflux is less relevant. A reward function handling such scenarios would require several “if-then-else” statements. These can be included in reward scores but are not always trivial to define and set-up for edge cases. In addition, multiparameter optimization becomes increasingly challenging when there are many constraints to fulfill [41]. In brief, the biggest challenge of deep generative methods is to define relevant reward scores, and this is unfortunately less studied.

Simple drug-likeness rules, multivariate methods for DMPK properties (solubility, permeability, clearance, etc.) and safety, as well as docking scores are typically included in reward scores as filters. However, several thousands of compounds will inevitably still pass those filters. This is related to the common lack of sufficient high-quality data, and the fact that we still often

struggle with making predictions to the required accuracy. Prediction of biological activity is an extremely hard problem since many phenomena involved are difficult to quantify precisely. Standard docking scores are most often not sufficient. Although, at times, methods such as Free Energy Perturbation (FEP) can improve the scoring accuracy for small perturbations of one structure into another, but not for major structural changes [42]. The use of FEP combined with active learning is gaining traction and is showing promise [43, 44]. Nonetheless, when the output contains many structurally diverse molecules, as frequently is the case for expressive SMILES-based generative methods, current methods' accuracies are not sufficient to filter down many compounds to a selected few. Despite the increasing prevalence of physics-based models in generative modelling, bioaffinity prediction remains very challenging.

Here, we propose shape and electrostatic potential matching as a strategy to bias generative models to propose compounds with different fragments (that are likely bioisosteres) of known lead compounds.

Using similarity as a design strategy

As mentioned above, current generative AI methods generally suffer from the lack of prediction accuracy. Thus, learning from past drug hunting experiences, we deliberately bias the AI method to generate compounds that are similar to active molecules already discovered. We approach this problem by relying on the similarity principle [45], which states that similar molecules tend to have similar properties [46]. Some advantages to this approach are discussed below. First, by generating molecules similar to the initial set available in the project, confidence in the predictions can be high because they remain in the applicability domain of the model. This is contrary to expressive methods that are designed to fully explore chemical space and generate structurally diverse compounds, which are consequently also the most uncertain to predict. Second, for similar compounds, the same chemical intermediates and established synthetic routes can often be re-used, facilitating speedy progress. Third, sometimes certain structural fragments (e.g., "privileged structures" [47]) are difficult to replace without severe drops in potency due to specific ligand-protein interactions.

As a related example, the strategy of molecular optimization using similarity was recently applied by Zhavoronkov and coworkers. They reported that deep learning enabled rapid identification of potent DDR1 kinase inhibitors [48]. Walters and Murcko analyzed the Zhavoronkov *et al.* study and reported that the AI-generated compound **B** (Figure 8) shared a common substructure with an already marketed multi-kinase inhibitor (Ponatinib, Figure 8), which was indeed included in the training set [49]. In some more detail, they ring-closed a benzamide carbonyl into an isoxazole moiety to yield an equipotent and unique compound [48]. These two compounds are very similar with regards to shape and electrostatics, see Figure 8. Thus, Zhavoronkov's AI method successfully mimicked

typical medicinal chemistry behavior, keeping certain parts fixed and making minor modifications to other.




Figure 8. Designing similar compounds can be a good tactic in drug discovery. Here illustrated with two potent DDR1 kinase inhibitors. The AI-generated compound **B** by Zhavoronkov *et al.* and ponatinib, a marketed multi-kinase inhibitor [48]. The compounds share a rather large common substructure. The hydrogen-bond acceptor and donor functionalities are visualized with electrostatic contours (red: negative, blue: positive). The ESP-Sim Tanimoto value is 0.81 for this pair.

It is sometimes believed that computer-aided design (CAD) methods need to provide radically "non-intuitive" different compounds to merit its use. However, believing that CAD approaches should surprise us and produce results that we would not have expected is a tall order. In this context, the scoring functions used in generative methods for reinforcement learning are not designed to extrapolate, and do not account for all aspects involved in designing drugs process. Thus, the power of current AI's lies more in pattern recognition than in creative discovery.

Over 100 different deep generative methods have been published the last couple of years. The methods are innovative and perform well in benchmark studies that measure models' ability to, for example, reproduce property distributions, and generate valid, diverse, and novel molecules [50]. One may thus conclude that generative modelling is essentially a solved problem – given a reward function, we now have the methods to generate molecules that satisfy it. Despite this success, biology and drug discovery remain immensely complex, and it is our viewpoint that current generative methods best serve to augment drug design. To take the next step (full autonomy), calculated predictions need ultra-high accuracy and for that we need to develop a broader understanding of human biology. The state of AI in drug design may be seen as analogous to the automotive industry. While the future of autonomous vehicles is promising and exciting, we are not near fully autonomous cars yet. Candidate drugs, as well as cars, still require human attention, given the complexity involved and the vast amount of edge cases that are non-trivial to code up efficiently. Thus, humans (with domain knowledge) are still very much needed in the process: to steer the tools and triage the output. In this context, we would like to highlight the Gruenif.ai tool where the user can provide feedback interactively while molecules are generated [51]. Such "human-in-the-loop" methods can be very effective. Future versions of DeepFMPO will include such functionalities.

CONCLUSIONS

The use of sophisticated computational methods for de novo design is attractive and deep generative methods have gained a lot of attention. Significant progress has been made when it comes to *generating* molecules. However, *scoring* them accurately remains a major challenge. Real-life project experience informs us that *in silico* predictions (e.g., synthesis, potency, properties) are constantly improving, but they are generally not accurate enough to prioritize a handful compounds for synthesis from a long list of high-scoring AI generated molecules. Thus, what really needs solving is being able to do ultra-accurate predictions to advance the field to the next level. Until then, the approach of biasing molecular design towards compounds similar to known actives will remain as one pragmatic and fruitful way to success.

Here we present a 3D fragment-based reinforcement learning approach for the generation of novel molecules with optimized properties, called “DeepFMPO v3D”. We furthermore developed a python package, ESP-Sim, for calculating molecular shape and electrostatic similarities. The inclusion of ESP scores promotes the generation of compounds similar to existing lead molecules, towards desirable sweet spot properties. The proposed method allows the calculation of state-of-the-art partial charges (e.g., RESP with B3LYP/6-31G**) obtained using the quantum chemistry program Psi4. In a scaffold-hopping case-study we show that our approach of using shape and electrostatics similarities performs well. DeepFMPO v3D ranks known equipotent scaffolds higher and generates them earlier (*i.e.*, speedier). The way the 3D method is implemented makes the approach essentially alignment-independent (on a molecular level) and does not require knowing the bioactive conformation. Both DeepFMPO v3D and ESP-Sim are freely available online.

AUTHOR INFORMATION

Corresponding Author

* Telephone number: +46 31 7065251. E-mail: jonas.bostrom@astrazeneca.com.

ACKNOWLEDGMENTS

The helpful code advice from Niclas Ståhl, and the constructive feedback on earlier versions of the manuscript from colleagues is gratefully acknowledged. E.H. acknowledges support from the Austrian Science Fund (FWF), project J4415, and the Machine Learning for Pharmaceutical Discovery and Synthesis Consortium (MLPDS).

SUPPLEMENTARY INFORMATION

Supplementary information associated with this article can be found online. The ESP-SIM code can be found here: <https://github.com/hesther/epsim> and the DeepFMPO_v3D

code version together with the associated data sets used in this study here: <https://github.com/giovanni-bolcato/deepFMPOv3D>.

AUTHOR CONTRIBUTIONS

E.H. developed and implemented the ESP-Sim method. G.B. implemented ESP-Sim into DeepFMPO and performed the machine learning experiments. The project was devised by J.B. All authors contributed to the writing of the manuscript.

REFERENCES

1. Hann, M. M.; Keserü, G. M. Finding the sweet spot: the role of nature and nurture in medicinal chemistry. *Nat. Rev. Drug Discov.* 2012, 11, 355–365.
2. Elton, D.C., Boukouvalas, Z., Fugea, M.D., Chunga, P. W. Deep learning for molecular design – a review of the state of the art *Mol. Syst. Des. Eng.* 2019, 4, 828–849
3. Merz Jr., K.M., De Fabritiis, G., Wei, G. JCI special issue on generative models for molecular design. *J. Chem. Inf. Model.* 2020, 60, 1072.
4. Sanchez-Lengeling, B., Aspuru-Guzik, A. Inverse molecular design using machine learning: Generative models for matter engineering. *Science* 2018, 361, 360–365.
5. Schneider, G., Fechner, U. Computer-based de novo design of drug-like molecules, *Nat. Rev. Drug Discov.* 2005, 4, 649–663
6. Gupta, A., Müller, A.T., Huisman, B.J.H., Fuchs, J.A., Schneider, P., Schneider, G. Generative recurrent networks for de novo drug design. *Mol Inform.* 2018, 37, 1700111
7. Popova, M., Isayev, O., Tropsha, A. Deep reinforcement learning for de novo drug design. *Sci. Adv.* 2018, 4, eaap7885
8. Weininger, D. SMILES, a chemical language and information system. 1. introduction to methodology and encoding rules. *J. Chem. Inf. Comp. Sci.* 1988, 28, 31–36.
9. Dalke, A. Weininger's Realization (accessed October 10, 2021) http://www.dalkescientific.com/writings/diary/archive/2016/12/02/Weiningers_realization.html
10. Merk, D.; Friedrich, L.; Grisoni, F.; Schneider, G. De novo design of bioactive small molecules by artificial intelligence, *Mol. Inf.* 2018, 37, 1700153

11. Ståhl, N., Falkman, G., Karlsson, A., Mathiason, G., Boström, J. Deep reinforcement learning for multiparameter optimization in de novo drug design *J. Chem. Inf. Model.* 2019, 59, 3166–3176.
12. S. Muchmore, A. J. Souers, I. Akritopoulou-Zanze, The use of three-dimensional shape and electrostatic similarity searching in the identification of a melanin-concentrating hormone receptor 1 antagonist *Chem. Biol. Drug Des.*, 2006, 67, 174–176.
13. Churchill, G.C. Identification of a chemical probe for NAADP by virtual screening *Nat. Chem. Biol.* 2009, 4, 220–226.
14. Boström, J. Grant, J.A.; Fjellström, O.; Thelin, A.; Gustafsson D. Potent fibrinolysis inhibitor discovered by shape and electrostatic complementarity to the drug tranexamic acid, *J. Med. Chem.* 2013, 56, 8, 3273–3280.
15. Skalic, M.; Jimenez, J.; Sabbadin, D.; De Fabritiis, G. Shape-based generative modeling for de-novo drug design. *J. Chem. Inf. Model.* 2019, 59, 1205–1214.
16. Imrie, F., Bradley, A. R., van der Schaar, M. Deane, C. M Deep generative models for 3D linker design *J. Chem. Inf. Model.* 2020, 60, 1983–1995.
17. Zhou, Z., Kearnes, S., Li, L. et al. Optimization of Molecules via Deep Reinforcement Learning. *Sci Rep* 9, 10752 (2019). <https://doi.org/10.1038/s41598-019-47148-x>
18. Fischmann, T.O.; Hruza, A.; Duca, J.S.; Ramanathan, L.; Mayhood, T.; Windsor, W.T.; Le, H.V.; Guzi, T.J.; Dwyer, M.P.; Paruch, K.; Doll, R.J.; Lees, E.; Parry, D.; Seghezzi, W.; Madison, V. Structure-guided discovery of cyclin-dependent kinase inhibitors. *Biopolymers*, 2008, 89, 372–379.
19. Riniker, S.; Landrum, G. A. Better informed distance geometry: Using what we know to improve conformation generation *J. Chem. Inf. Comp. Sci.* 55:2562-74 (2015).
20. Landrum, G.A. RDKit: open-source cheminformatics software, version 2021.03.1 <http://www.rdkit.org>
21. Hare, B.J.; Walters, W.P.; Caron, P.J.; Bemis, G.W. CORES: An automated method for generating three-dimensional models of protein/ligand complexes. *J. Med. Chem.* 2004, 47, 19, 4731–4740.
22. Boström, J., Hogner, A. Schmitt, S. Do structurally similar molecules bind in a similar fashion *J. Med. Chem.* 2006, 49, 6716–6725.
23. Tanimoto, T.T. An elementary mathematical theory of classification and prediction, IBM Report (November, 1958), cited in: G. Salton, Automatic Information Organization and Retrieval (McGraw-Hill, 1968) p. 238.
24. Willett, A. P.; Barnard J.M.; Downs, G.M. Chemical similarity searching, *J. Chem. Inf. Comput. Sci.* 1998, 38, 983-996.
25. Carbó, R.; Leyda, L.; Arnau, M. How similar is a molecule to another? An electron density measure of similarity between two molecular structures. *Int. J. Quantum Chem.* 1980, 17, 1185-1189.
26. Carbó, R.; Domingo, L. LCAO-MO similarity measures and taxonomy. *Int. J. Quantum Chem.* 1987, 32 517-545.
27. Good, A. C.; Hodgkin, E. E.; Richards, W. G. (1992). Utilization of Gaussian functions for the rapid evaluation of molecular similarity. *J. Chem. Inf. Comput. Sci.* 1992, 32, 188-191.
28. Turney, J. M.; Simmonett, A. C.; Parrish, R. M.; Hohenstein, E. G.; Evangelista, F. A.; Fermann, J. T.; Mintz, B. J.; Burns, L. A.; Wilke, J. J.; Abrams, M. L.; Russ, N. J.; Leininger, M. L.; Janssen, C. L.; Seidl, E. T.; Allen, W. D.; Schaefer, H. F.; King, R. A.; Valeev, E. F.; Sherrill, C. D.; Crawford, T. D. Psi4: an open-source ab initio electronic structure program *WIREs Comput. Mol. Sci.* 2012, 2, 556–565.
29. The RESP code is a plug-in to the Psi4 quantum chemistry package available from <https://github.com/cdsgroup/resp> (Sherrill Research Group).
30. Dask is a flexible library for parallel computing in Python. <https://docs.dask.org/en/latest/> (accessed June 22, 2021)
31. GAMESS parameters <https://github.com/streaver91/gamess/blob/master/source/prplib.src>
32. Singh, U.C.; Kollman P.A. An approach to computing electrostatic charges for molecules. *J. Comp. Chem.* 1984, 5, 129–145.
33. Ertl P. In silico identification of bioisosteric functional groups. *Curr. Opin. Drug Discov. Develop.* 2007, 10, 281–288.

34. Meanwell NA. Synopsis of some recent tactical application of bioisosteres in drug design. *J. Med. Chem.* 54, 2529–2591 (2011).
35. Falaguera M.J.; Mestres J. Congenericity of claimed compounds in patent applications. *Molecules* 2021, 26, 5253.
36. Bento, A.P.; Gaulton, A.; Hersey, A.; Bellis, L.J.; Chambers, J.; Davies, M.; Krüger, F.A.; Light, Y.; Mak, L.; McGlinchey, S.; Nowotka, M.; Papadatos, G.; Santos, R.; Overington, J.P. The ChEMBL bioactivity database: an update (version 26) *Nucleic Acids Res.*, 42 (2014), pp. D1083-D1090.
37. Wildman, A.; Crippen, G. M. *J. Chem. Inf. Comput. Sci.* 199, 39, 868-873.
38. Meyers, J.; Fabian, J.B.; Brown, N. De novo molecular design and generative models, *Drug Disc. Today*, 2021, in press, <https://doi.org/10.1016/j.drudis.2021.05.019>.
39. Brenner, S. Sequences and consequences. *Phil. Trans. R. Soc. B.*, 2010, 365, 207–212.
40. Gao, W., Coley, C.W. The synthesizability of molecules proposed by generative models *J. Chem. Inf. Model.* 2020, 60, 12, 5714–5723.
41. Jin, W., Barzilay, R., Jaakkola, T. Composing molecules with multiple property constraint <https://arxiv.org/abs/2002.03244>
42. Wang, L., Wu, Y., Deng, Y., Kim, B., Pierce, L., Krilov, G., Lupyan, D., Robinson, S., Dahlgren, M.K., Greenwood, J. Accurate and reliable prediction of relative ligand binding potency in prospective drug discovery by way of a modern free-energy calculation protocol and force field. *J. Am. Chem. Soc.*, 2015, 137, 2695–2703.
43. Konze, K.D.; Bos, P.H.; Dahlgren, M.K.; Leswing, K.; Tubert-Brohman, I.; Bortolato, A.; Robbason, B.; Abel, R.; Bhat, S. Reaction-based enumeration, active learning, and free energy calculations to rapidly explore synthetically tractable chemical space and optimize potency of cyclin-dependent kinase 2 inhibitors *J. Chem. Inf. Model.* 2019, 59, 3782-3793.
44. Ghanakota P, Bos P, Konze K, Staker J, Marques G, Marshall K, et al. Combining cloud-based free energy calculations, synthetically aware enumerations and goal-directed generative machine learning for rapid large-scale chemical exploration and optimization. ChemRxiv. <https://doi.org/10.26434/chemrxiv.11825679.v1>
45. Johnson, A. M.; Maggiora, G. M. (1990). *Concepts and Applications of Molecular Similarity*. New York: John Wiley & Sons. ISBN 978-0-471-62175-1.
46. Martin, Y. Do structurally similar molecules have similar biological activity? *J. Med. Chem.* 2002, 45, 19, 4350-4358.
47. Evans, B.E.; Rittle, K.E.; Bock, M.G.; DiPardo, R.M. et al. Methods for drug discovery: development of potent, selective, orally effective cholecystokinin antagonists. *J. Med. Chem.* 1988, 31, 2235-224.
48. Zhavoronkov, A.; Ivanenkov, Y.A.; Aliper, A. et al. Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nat. Biotechnol.* 2019, 37, 1038–1040.
49. Walters, P., Murcko, M. Assessing the impact of generative AI on medicinal chemistry *Nat. Biotechnol.* 2020, 38, 143–145.
50. Brown, N., Marco Fiscato, M., Segler, M.H.S., Vaucher, A. C. GuacaMol: benchmarking models for de novo molecular design. *J. Chem. Inf. Model.* 2019, 59, 3, 1096-1108.
51. Winter, R., Retel, J., Noe, F., Clevert, D.-A., Steffen, A. Grünifai: interactive multiparameter optimization of molecules in a continuous vector space. *Bioinformatics*, 36 (2020), 4093-4094.

