

Deep Generative Models for Ligand-based de Novo Design Applied to Multi-parametric Optimization

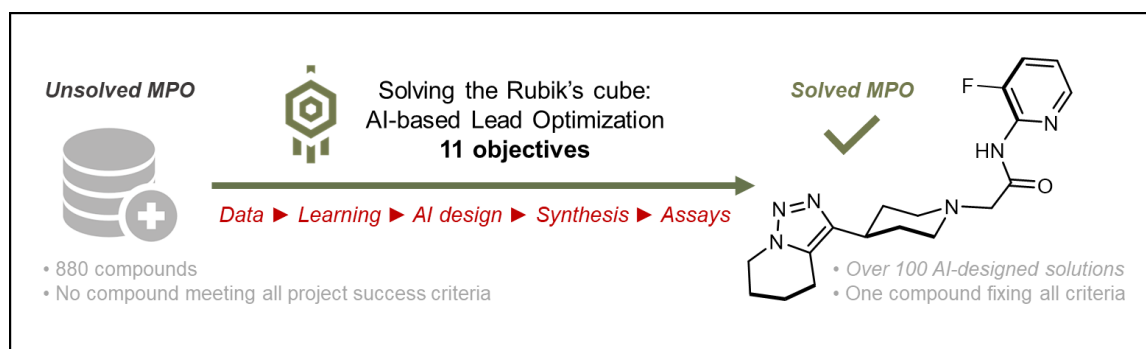
Quentin Perron ⁽¹⁾, Olivier Mirguet ⁽²⁾, Hamza Tajmouati ⁽¹⁾, Adam Skiredj ⁽¹⁾, Anne Rojas ⁽²⁾, Arnaud Gohier ⁽²⁾, Pierre Ducrot ⁽²⁾, Marie-Pierre Bourguignon ⁽²⁾, Patricia Sansilvestri-Morel ⁽²⁾, Nicolas Do Huu ⁽¹⁾, Françoise Gellibert ⁽²⁾, Yann Gaston-Mathé ^{(1)*}

⁽¹⁾ Iktos, 65 rue de Prony, 75017 Paris, France

⁽²⁾ Institut De Recherches Servier, 125 Chemin de Ronde, 78290 Croissy, France & 11 rue des Moulineaux, 92150 Suresnes, France

*e-mail : yann.gaston.mathe@iktos.com

TABLE OF CONTENTS



ABSTRACT

Multi-Parameter Optimization (MPO) is a major challenge in New Chemical Entity (NCE) drug discovery projects, and the inability to identify molecules meeting all the criteria of lead optimization (LO) is an important cause of NCE project failure. Several ligand- and structure-based de novo design methods have been published over the past decades, some of which have proved useful multiobjective optimization. However, there is still need for improvement to better address the chemical feasibility of generated compounds as well as increasing the explored chemical space while tackling the MPO challenge. Recently, promising results have been reported for deep learning generative models applied to de novo molecular design, but until now, to our knowledge, no report has been made of the value of this new technology for addressing MPO in an actual drug discovery project. Our objective in this study was to evaluate the potential of a ligand-based de novo design technology using deep learning generative models to accelerate the discovery of an optimized lead compound meeting all in vitro late stage LO criteria.

Key words: Artificial Intelligence, Multiparameter Optimization, Drug Discovery, Lead Optimization

1. INTRODUCTION

Drug design is a challenging task. From hit identification to hit-to-lead and lead optimization, the quest to discover a new chemical entity (NCE) with desired properties is burdensome. Exploration of a nearly infinite chemical space (10^{60} drug-like molecules is a low range figure)^{1,2} is required in order to solve a multi-parametric optimization (MPO) challenge: identifying the rare compounds which satisfy a blueprint encompassing criteria such as biological activity, selectivity, (lack of) toxicity, pharmacokinetics (*i.e.* DMPK), synthetic accessibility and finally novelty.^{3,4} The average cost to develop a preclinically validated drug candidate is estimated around \$50 million, and drug design, more specifically lead optimization, represents the lion's share (~70%) of the cost of preclinical research.⁵

Structure- and ligand-based computer aided drug design (CADD) technologies (*e.g.* docking, QSAR, etc.), which have been developed to improve the productivity of the drug design process, have brought notable progress over the last decades.^{6,7} Still, most classical CADD approaches have focused on the prediction of molecular properties rather than on the exploration of the chemical space to identify novel compounds with optimal properties. Such *in silico* exploration of the chemical space has mostly been performed through the virtual screening of pre-existing or virtual compound libraries, with the exploration being intrinsically restricted to the initial compound library.^{8,9} Graph-based genetic algorithms, sometimes used for *in silico* chemical optimization, have had limited success and are mostly limited by transformation rules.¹⁰ More recently, the development of artificial intelligence (AI) approaches in CADD, and more specifically *de novo* drug design through the use of deep generative models, has triggered a lot of interest in the CADD community.¹¹ Indeed, in contrast to traditional virtual screening methods, deep generative models allow to effectively explore the chemical space by designing new molecules and optimizing the desired parameters during the generative process.

Generative models for molecular design can be characterized by three main features: which molecular representation they use, how they generate molecules, and how they perform property optimization. Many methods have been reported, each with different approaches regarding those features: *i)* The molecular representation can be either text (SMILES,^{12,13} SELFIES,^{14,15} DeepSMILES¹⁶), or a graph or a set of fragments.¹⁷⁻¹⁹ *ii)* The molecule generation strategy can use a simple policy, for instance: add or remove atoms or bonds.¹⁷ It can also rely on deep generative models such as recurrent neural networks (RNNs), auto-encoders (AEs) or generative adversarial networks (GANs).^{20,21} *iii)* The property optimization strategy can be based on reinforcement learning,^{17,22-23} continuous optimization,²⁰ Bayesian optimization,²⁴ genetic algorithms¹⁵ or particle swarm optimization.²⁵

Despite the amount of research in generative modeling and its potential to allow an efficient exploration of the chemical spaces to identify new molecules with the desired *in silico* properties, evidence of the benefit of such AI-based approaches to solve MPO issues in complex real-life cases is still elusive, and AI-based drug design is perceived as overhyped by a significant part of the chemists' and chemo-informaticians' community.²⁶

As previously stated, MPO is a major challenge in NCE drug discovery projects, and the inability to identify molecules meeting the Target Product Profile (TPP) in LO is an important cause of NCE project failure or delay. Herein, we describe the application of a ligand-based *de novo* design AI technology based on deep generative models in a real-life LO stage drug discovery project and its impact on fostering the discovery of optimized lead compounds meeting the project's TPP criteria.

2. MATERIALS AND METHODS

2.1. Project dataset

The selected project was a live internal drug discovery project at LO stage that had been running for several years. The project dataset consisted in a library of 881 molecules with associated bioactivity measurements from 11 biological assays: one primary activity assay (undisclosed phenotypic assay: % of activation at 30nM), 6 off-target activity assays (selectivity criteria on 5-HT2A, 5-HT2B, α 1, D1, Na_v1.2, hERG: % of inhibition) and 4 ADME assays (microsomal stability on human (HLM) and rat (RLM): % of stability; permeability and efflux Caco2 assays: % of absorption and efflux ratio). For each objective, a threshold value was defined according to the target product profile (TPP) designed by the project team. A summary of the thresholds, percentage of compounds measured and percentage of compounds meeting the required threshold for each assay is reported in table 1.

Objectives	Activity	5-HT2A	5-HT2B	α 1	D1	Na _v 1.2	hERG	RLM	HLM	Caco-2 FAbs	Caco-2 Efflux
Concentration	30 nM	10 μ M	10 μ M	10 μ M	10 μ M	10 μ M	10 μ M	-	-	-	-
Filled %	29%	28%	26%	33%	28%	30%	59%	90%	90%	87%	77%
Blueprint Threshold	$\geq 30\%$	$\leq 50\%$	$\leq 50\%$	$\leq 50\%$	$\leq 50\%$	$\leq 50\%$	$\leq 30\%$	$\geq 50\%$	$\geq 50\%$	$\geq 90\%$	≤ 15
In blueprint rate	59%	29%	35%	33%	53%	68%	45%	49%	35%	61%	80%

Table 1: Statistical outlook of the initial dataset (Each column represents an assay and the concentration at which compounds were tested. "Filled %" describes the % of molecules in the dataset which have data in the assay. "Blueprint Threshold" is the value set as the objective to achieve in each assay. "In blueprint rate" is the percentage of molecules meeting each objective individually.)

The dataset was sparse, with 10-70% missing data rates depending on the objectives. Because of the specificity of the primary assay, a complex *ex vivo* phenotypic assay, the ADME assays were very well documented, whereas only 251 compounds had been measured in the primary activity and selectivity assays. The dataset was well balanced, with >50% compounds meeting individually most objectives, with lower rates (~30-35%) observed for 5-HT2A, 5-HT2B, α 1 and HLM.

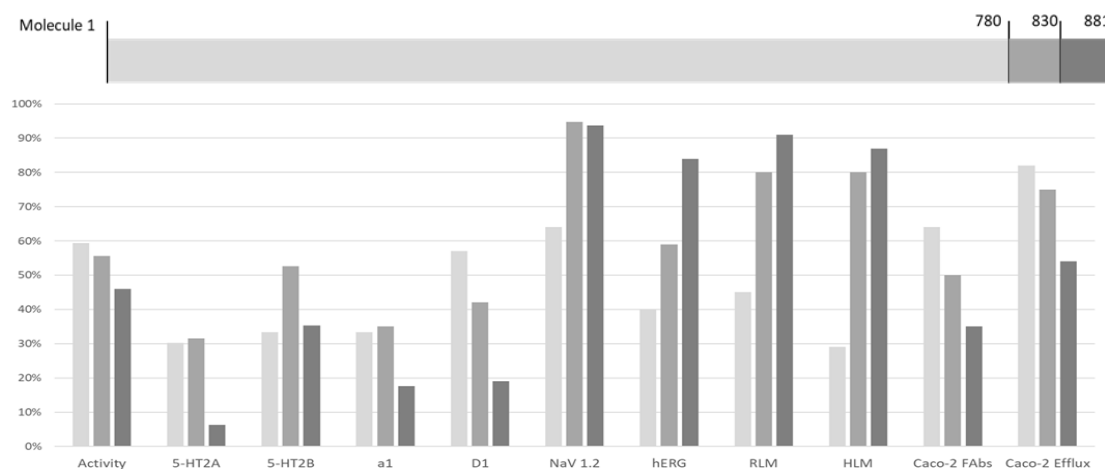


Figure 1: % of molecules in the initial project dataset meeting the different objectives along the chronology of the project (light grey: molecule 1 to 780; medium grey: molecule 781 to 830; dark grey: molecule 831 to 881).

The evolution of the percentage of compounds meeting each objective during the chronology of the project is displayed in figure 1. It shows that the project team has been able to substantially increase the performance across iterations for Na_v1.2, hERG, RLM, HLM, with 80-90% of designed molecules meeting the required goal at the end of the program. Conversely, performance had strikingly dropped on 5-HT_{2A}, α 1, D1, and permeability assays. As an example, only 6% of the last fifty molecules synthesized met the 5-HT_{2A} selectivity objective.

In the subset of 48 molecules out of 881 which had been measured against all 11 objectives, the average number of objectives met was 6.4 out of 11. Among these, 6 molecules appeared to have a promising profile, meeting 9 objectives out of 11 (table 2). Molecule 732 (**mol 732**) depicted in **Figure 2** was the best compound in the whole dataset, meeting all objectives except absorption, which was nearly met, and efflux.

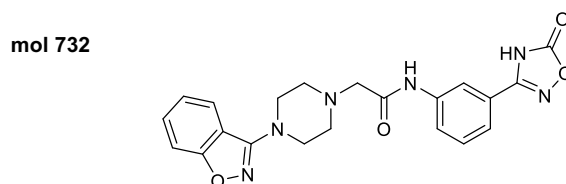


Figure 2: Structure of the best molecule, **mol 732**, of the initial dataset.

It is worth noting that the 1,2-benzisoxazole in **mol 732** was also found in 61% of the project's compounds, and in 78% of the last 50 compounds made by the project team, indicating the importance that had been given by the medicinal chemistry team to that substructure, as a seemingly promising avenue for achieving a good balance between all desired properties. Only a couple of piperidine and piperazine linkers were used throughout the project, while more variability had been introduced in the East part heterocycles.

Activity	5-HT _{2A}	5-HT _{2B}	α 1	D1	Na _v 1.2	hERG	RLM	HLM	Caco-2 FABs	Caco-2 Efflux	
194.0	20.0	18.0	1.0	4.0	0.0	19.0	82.85	63.35	88.99	26.2	mol 732
83.0	69.0	-25.0	45.0	6.0	13.0	6.4	69.04	31.93	97.6	1.96	mol 663
46.0	46.0	69.0	14.0	14.0	-14.0	25.8	60.28	25.43	98.86	0.75	mol 559
48.0	71.0	48.0	12.0	14.0	39.0	25.0	68.83	33.58	99.37	0.39	mol 555
115.0	76.0	15.0	37.0	-3.0	-13.0	5.4	80.82	83.54	72.24	12.3	mol 550
46.0	6.0	44.0	29.0	-11.0	20.0	12.4	93.11	78.36	73.8	34.1	mol 435

Table 2: Biological profiles of the most promising lead molecules in the initial project dataset.

Also worthy of note, as shown in Table 2, a compound with a promising profile, **mol 435**, meeting 9 out of 11 objectives but missing absorption and efflux, quite close to **mol 732** in terms of biological profile, had been obtained much earlier in the project. 297 additional molecules were needed to partially improve the overall compound profile. During the design process from **mol 435** to **mol 732**, permeability objectives were met in three molecules (**mol 555**, **mol 559** and **mol 663**), but only to the detriment of 5-HT_{2A/B} selectivity or metabolic stability.

2.2. QSAR models development

Bioactivity data were binned according to TPP thresholds (*i.e.*, 1 if meeting the TPP specification, else 0). Eleven independent QSAR models were developed using ridge logistic regression based on Morgan fingerprint molecular representations.²⁷ The Morgan fingerprint was built without including chirality (two stereoisomers have an identical fingerprint) as most of the molecules in the dataset were achiral and the stereochemistry was known. No specific processing of tautomers was performed (two different tautomers of the same molecule have different fingerprints and likely different scores).

Model selection was performed using k-fold (k=4) cross-validation. It concerned two parameters: the penalty parameter and the operating threshold probability. The penalty parameter was selected to maximize the ROC AUC.²⁸ Once the penalty parameter selected, the operating threshold probability to predict conformity to TPP (noted as 1 in figure S1) was selected on the former k-folds to maximize precision to the detriment of recall, in order to reduce the risk of false positives. The best model, trained on 80% of the data (*i.e.*, training set) was subsequently tested on the remaining 20% of the initial dataset (*i.e.*, test set).

On average, the QSAR predictive models performed well with high precision in the test sets, except for 5-HT2B (precision 67%). Interpretability of the results was difficult for activity, alpha 1 and 5-HT2A due to the small number of positive compounds in the test set (confusion matrices are provided in figure S1, ROC AUC plots are provided in figure S2). The selected models were then trained on the whole dataset before switching to the generative phase of our work.

2.3. Generative model

As explained above, many architectures of molecular deep generative models have been reported to date. At the time this study was conducted (it was initiated in 2017), fewer architectures had been published. Molecule generation and property optimization strategies were inspired by Segler *et al.* which uses a deep recurrent neural network (RNN) generator.¹³

2.3.1. Molecule generation strategy

A deep recurrent neuronal network (RNN), and more precisely a deep Long Short-Term Memory (LSTM), was used to generate molecules represented as SMILES.^{12,29} The LSTM was first trained on the ChEMBL database, using teacher forcing,³⁰ to build a character-based language model for generating SMILES strings.¹³

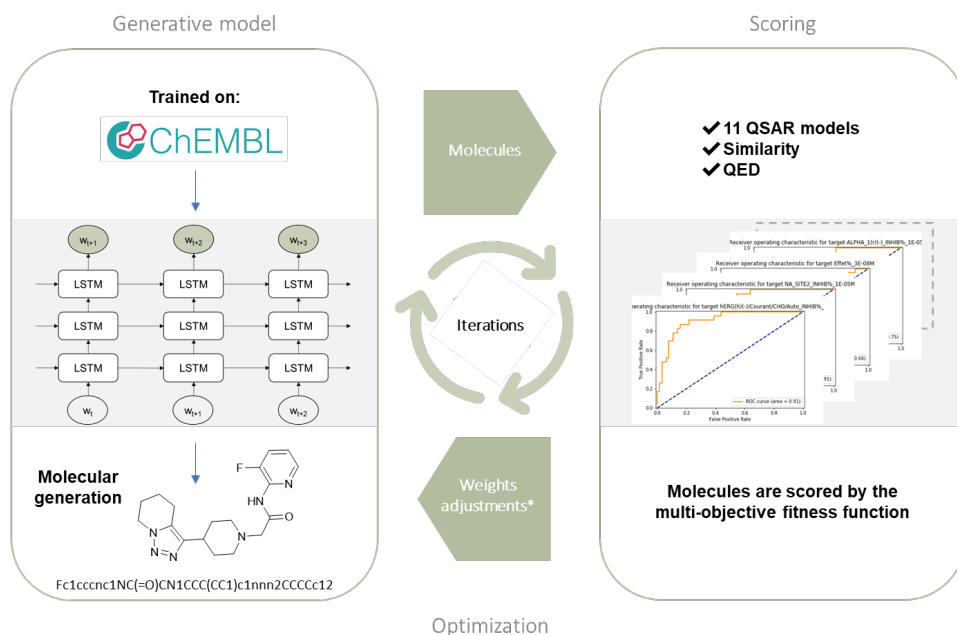
It is reminded that the role of a language model p is to model the next character probability distribution given the sequence of previous characters:

$$p(x_{t+1}|x_1x_2 \dots x_t) = LSTM(x_{t+1}|x_1x_2 \dots x_t)$$

SMILES are generated by iteratively sampling the next character from its inferred past conditioned distribution $p(x_{t+1}|x_1x_2 \dots x_t)$. Generating a SMILES starts and ends respectively with the special tokens of the vocabulary 'START' and 'END'.

The SMILES in the ChEMBL database were transformed into their canonical achiral RDKit version. No data augmentation by enumerating the different ways of writing a SMILES, nor by enumerating the tautomeric forms of the same compound was performed. Thus trained, the LSTM language model generates achiral SMILES. Identical compounds can be generated with different writings of their SMILES. Tautomers of the same compound are generated as distinct molecules.

After being trained on ChEMBL, the LSTM language model had a 94% SMILES chemical validity rate.



Scheme 1: Generative model architecture

2.3.2. Project dataset distribution learning

The LSTM trained on ChEMBL database has learnt to generate molecules belonging to ChEMBL chemical space.

In order to be scored, generated molecules should stay near the applicability domain of the QSAR models. This applicability domain can be approximated by the structural similarity to the project dataset.

Thus, the previous LSTM model was re-trained in teacher forcing on the project dataset. This second training allows to zoom in the chemical space studied so that QSAR models can be applied.

2.3.4. Molecule optimization strategy

The molecule optimization strategy that was used is named “Hillclimb-MLE”.^{13,23} It is an iterative process where the LSTM generative model is fine tuned in teacher forcing on an optimal set of SMILES that evolves over time as follow: step after step, this set of SMILES is updated by retaining only the top 10% of compounds generated since the first step. The optimality ranking was established using a scalar reward function that combines thirteen targets:

+ Eleven probabilities of activity (p_i) _{$1 \leq i \leq 11$} returned by the classifiers described above (QSAR models built on the training data set):

+ Similarity to the project dataset D , computed as:

$$S(mol) = \max\{Tanimoto_Similarity(mol, mol_j); mol_j \in D\}$$

+ QED.³¹

Denoting $(x_i)_{1 \leq i \leq 13}$ and $(T_i)_{1 \leq i \leq 13}$ respectively the vector of our thirteen targets of interest and their thresholds for being in the blueprint, the reward function used in this project was the following:

$$Reward((x_i)_{1 \leq i \leq 13}) = - \sum_{i=1}^{13} \log \left(\frac{x_i}{T_i} \right)$$

Thresholds of QSAR scores (*i.e.*, $(T_i)_{1 \leq i \leq 11}$) are their corresponding probability operating thresholds. The selected thresholds for similarity to the project dataset (*i.e.*, T_{12}) and QED (*i.e.*, T_{13}) are 0.5 and 0.4, respectively.

Results are seed dependent and thus, many runs were conducted, each one leading to new propositions to solve the problem.

2.4. Assessment and ranking of generated compounds

The generative algorithm designed 150 virtual compounds predicted to be optimal with regards to the project's TPP (*i.e.* predicted to meet the required threshold for all targets), and with reasonable complexity as assessed by a chemist (at the time of the study, no satisfying synthetic accessibility scoring tool was available to help prioritize compounds). Virtual candidates were ranked on their overall probability of being in the TPP, their drug likeness (QED),³¹ and their similarity to the initial dataset (*i.e.*, Tanimoto distance). Indeed, the applicability domain of the QSAR models is a critical point and must be carefully monitored to avoid false positives.

To help chemists assess the novelty and risk associated with the proposed molecules, a specific visualization was developed, by adapting the similarity map visualization.³² This visualization, which we have named "applicability map" (figure 3), enables to highlight, for each proposed molecule, the atoms which are either present or absent in the initial dataset, as follows: a) in green, the atoms which are very well known because they appear very often in the same chemical environment in the initial dataset (*i.e.*, the lead scaffold for instance); b) in red, the new atoms or atoms already known but appearing in a new position; and c) not highlighted: the atoms which have been seen before in the same position, but only a few times.

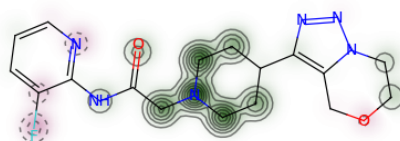


Figure 3: Applicability map of **mol 887**. The central core of the molecule is present in most of the compounds of the initial dataset and therefore very well known, the transparent parts indicate that atoms in those positions are existing in the initial dataset but only in few cases and 3 atoms are glowing in red (*i.e.*, the fluoropyridine moiety) indicating that this function has never been tried before.

The applicability map visualization helps to better detect potential false positives, since a molecule predicted to be active but with functional groups appearing as predominantly red in this representation indicates a higher level of risk associated with the prediction. The best compounds are molecules with green and transparent atoms, indicating that at the same time the structure is within the knowledge of the initial dataset (green part) and in a domain of the chemical space which is worth to be explored (transparent part). This representation facilitates the selection of compounds for

synthesis and test among the proposed AI-designed structures, which is a trade-off between risk and novelty, exploitation of the current dataset and exploration of new ideas suggested by the AI algorithm.

2.5. Compound selection & synthesis

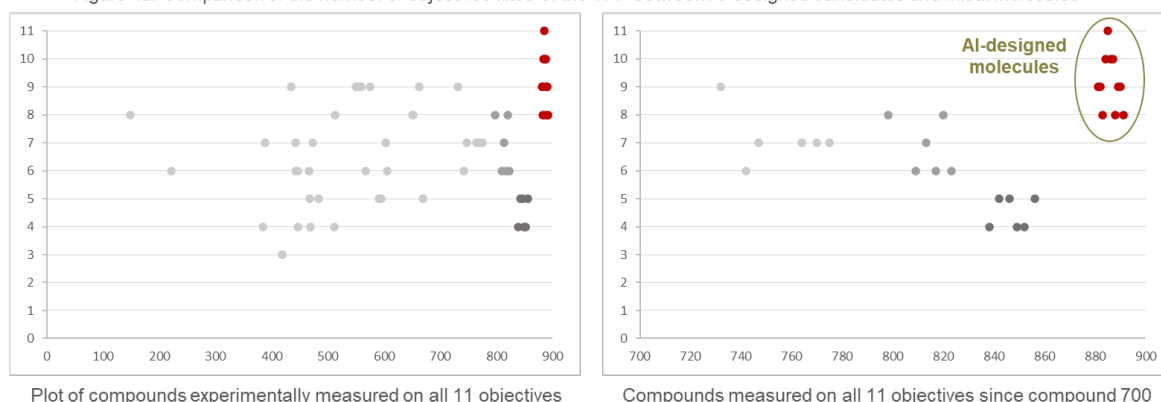
From the newly generated library, 20 AI-designed molecules out of the 150 were selected for synthesis and test based on their algorithmic ranking, structural novelty, synthetic accessibility, and consistency of the ADME predictions with those provided by global predictive models available at Servier. In the 3-week timeframe allowed for compound synthesis: 11 compounds were successfully synthesized and tested on all the project's assays (see figure 4, figure 5 and figure 6) whereas 9 molecules failed to be synthesized.

3. RESULTS

After synthesis and test, the AI-generated candidates were found to outperform the initial library, including the last 50 compounds made within the project. The average number of objectives met by the AI-designed compounds was 9.5 (*i.e.*, 86% success rate) *versus* 6.4 (*i.e.*, 58% success rate) previously. Moreover, the AI-generated molecules reversed the decreasing trend in TPP conformity observed in the last molecules of the library (figure 4a). Analysis shown in figure 4b illustrates that, compared to the initial dataset, novel molecules were better on activity (*i.e.*, in the blueprint 65% of the time) and excellent for all selectivity and permeability criteria (*i.e.*, over 90% of the time in the blueprint). Metabolic stability, however, was lower, with a 55% conformity rate. More importantly, from the 11 new compounds, one met simultaneously all 11 objectives of the TPP (figure 5) and two compounds met 10/11 objectives (figure 5 and 6), while being just below the required threshold, within the error margin of the assay, for the missed objective.



Figure 4a: Comparison of the number of objectives fixed of the TPP between AI-designed candidates and initial molecules



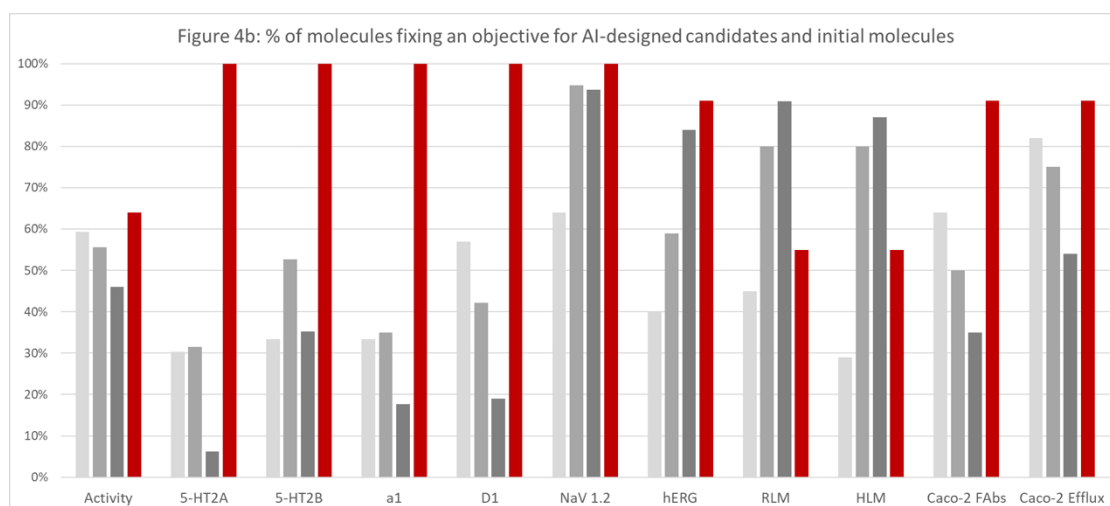


Figure 4: **a)** Number of objectives satisfied according to project's chronology (vertical axis: number of objectives satisfied / horizontal axis: chronologic numbering of compound / please note initial data was sparse with only 48 compounds tested on all criteria. **b)** Hit rate comparison between AI-designed candidates and initial molecules for each TPP objective.

The best AI-designed compound (**mol 885**), meeting all objectives, is represented in Figure 5. Notably, this compound contains a [1,2,3]triazolo[1,5-a]piperidine moiety which was very rare in the initial data set, appearing in only six molecules, and always correlated to poor permeability and efflux, which had led the project team to stop investigating this motif. It is remarkable that the AI algorithm retained that substructure, combining it with a 3-fluoropyridine in the East part, which had never been tried before. Surprisingly, the association of this discarded substructure with an unexplored heterocycle turned out to be a winning combination for solving the MPO objective of the project.

As a matter of fact, the 11 AI-designed compounds that were synthesized and tested displayed functional groups that were either rare in the initial dataset or never tried earlier in the project (see figure 6). It suggests that this method can propose significant innovations, by its ability to identify favourable modifications, even with few data to learn from.

Objectives	Activity	5-HT2A	5-HT2B	α1	D1	Na _v 1.2	hERG	RLM	HLM	Caco-2 FAbs	Caco-2 Efflux
Best AI designed compound	83	7	18	7	-9	2	3	57	75	97	7

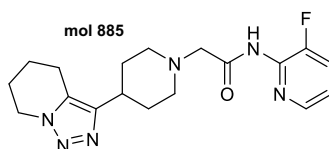


Figure 5: MPO profile of the best AI-designed molecule, **mol 885**.

One striking example is **mol 886**, where an aliphatic group was introduced in replacement of an aryl moiety, where only aromatic moieties had been used before at this specific position.

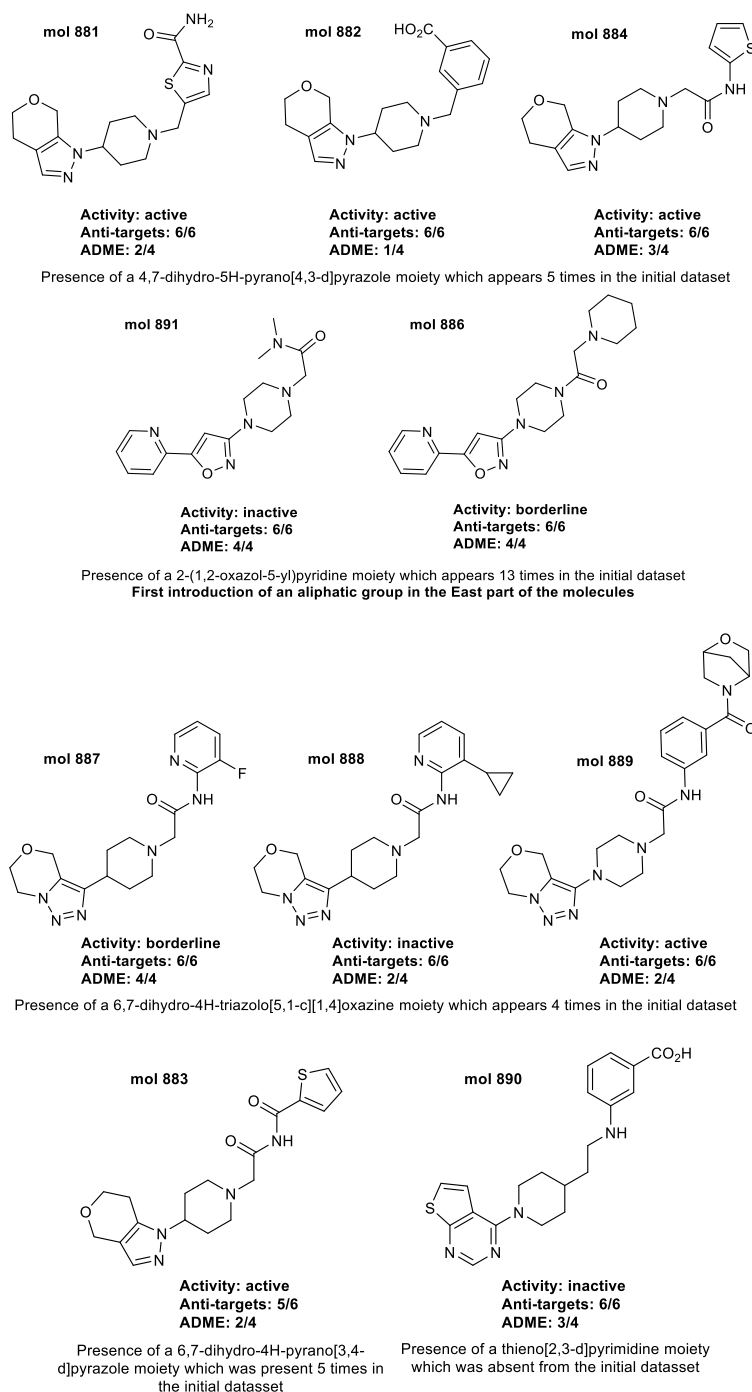


Figure 6: Structures and biological features of original compounds sampled by the DL algorithm. Borderline compounds have been measured below the desired activity threshold but within the error margin of the assay while active and inactive molecules were measured respectively above and below the threshold.

The AI algorithm was also able to optimize ADME properties in sub-series with specific ADME issues. For example, it was able to design permeable compounds within the 6,7-dihydro-4H-triazolo[5,1-c][1,4]oxazine sub-series while maintaining safety and stability, when all compounds in that sub-series had permeability issues. Likewise, within the pyrido-isoxazole series, compounds with reduced efflux were identified while maintaining safety and stability (figure 7).

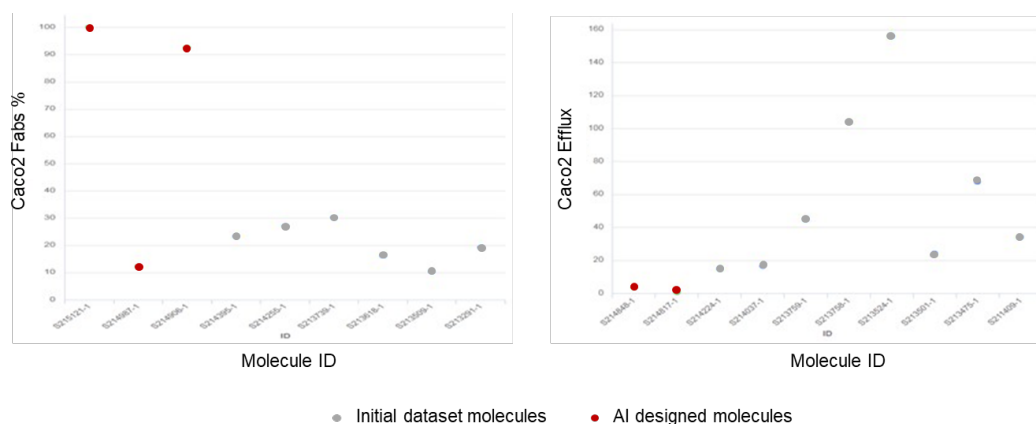


Figure 7: Permeability (left) and Efflux (right) properties of AI-designed vs original dataset compounds in the a) 6,7-dihydro-4H-triazolo[5,1-c][1,4]oxazine series (left) and b) pyrido-isoxazole series (right).

An analysis of the drug-likeness profile of the compounds based on their property forecast index (PFI),³³ molecular weight (MW) and sp³ fraction was performed (figure 8). Ten out of eleven AI-designed compounds were found to have a very favourable profile with low PFI, low MW, and high sp³ fraction, compared to the molecules from the initial data set.

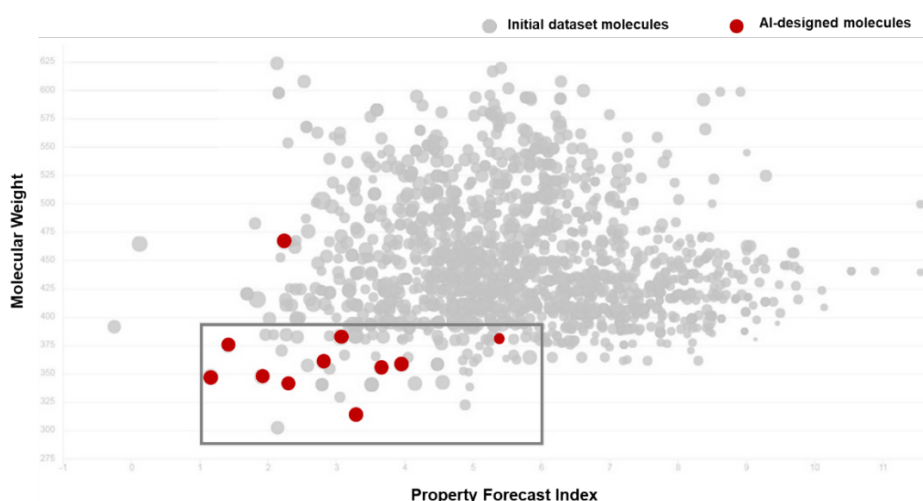


Figure 8: Plot of MW in function of PFI for initial molecules and AI-designed compounds

To provide insights about structural diversity and chemical space features of both the initial dataset and AI designed compounds, a principal component analysis (PCA) was computed on the Morgan fingerprints (*i.e.*, extended connectivity fingerprints (ECFP) of 1024 bits, radius 2) of the molecules in the dataset.³⁴ First, a representation of the 251 compounds from the initial library that were measured in the primary activity assay is provided (figure 9). This plot reveals the absence of a probability gradient or narrow area of activity since active molecules can be found in all areas of the explored chemical space. Conversely, a display of the number of objectives met by these 251 molecules allows to delineate an area where the MPO score is the highest (*i.e.* the upper left corner of the plot in figure 9).

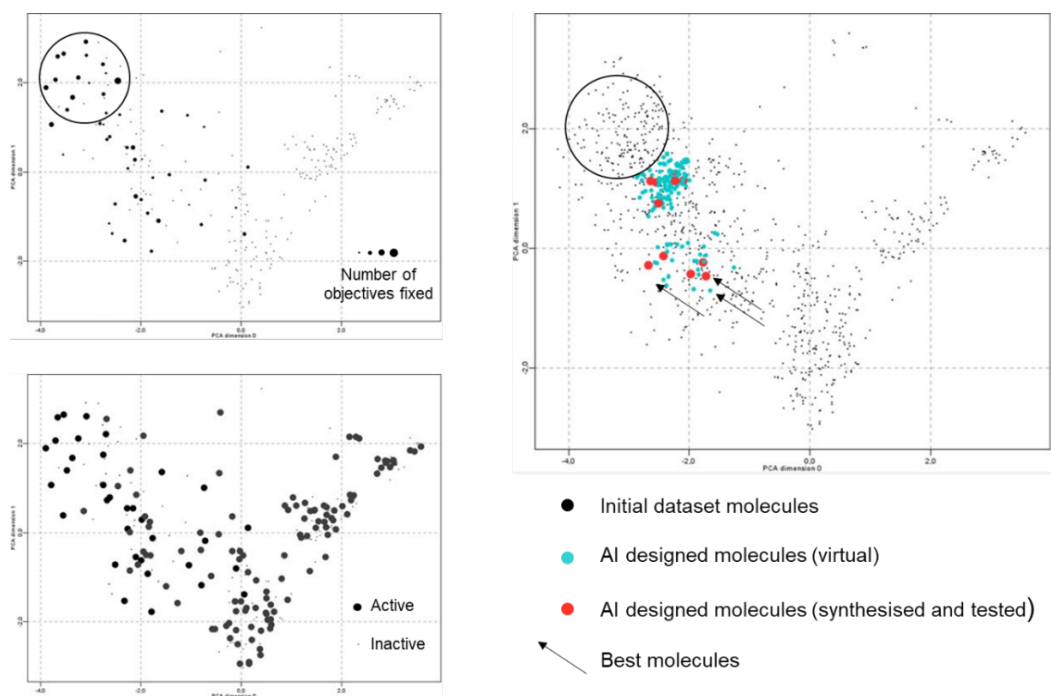


Figure 9: Bottom left) PCA of the 251 active compounds from the initial dataset. **Top left)** PCA of the same 251 active compounds with correlation to the TPP criteria hit rates. **Right)** Plot of the AI-designed molecules.

Strikingly (Figure 9), the AI algorithm did not design any molecule in that seemingly promising chemical space. All AI-designed structures are indeed located in a distinct yet specific area, demonstrating the capacity of this algorithm to come up with non-trivial solutions.

5. DISCUSSION

A typical hurdle of MPO is that optimization of some objectives leads to drop of performance in others, but the present method allowed to design compounds that were simultaneously optimized on eleven parameters. Yet, several features of the initial dataset were key to enable achieving such performance.

Overall, the performances of the models built to predict bioactivity on each assay were good, thereby validating the approach of project data-guided optimization. This requires enough data to build a decent model (in our case, the least documented assay had ~250 data points) and a reasonably well-balanced data set with enough compounds meeting each objective individually. Also, the generative model was able to find theoretical solutions to the MPO challenge within the chemical space of the project, meaning that based on the available data, there was indeed ways to solve the apparent anti-correlations between the objectives.

This favourable configuration may not be present in all cases, and the potential of the method to solve MPO challenges in more complex cases remains to be demonstrated. Several approaches could be envisaged to circumvent the lack of balanced data on some objectives, such as using generic models trained on large and diverse legacy data, e.g., for ADMET properties prediction, or using structure-based modelling to guide optimization on target or anti-targets if such structural information is available. To address the tricky issue of the inability to identify structures solving the MPO challenge

within the project's chemical space, adding an active learning component to the SAR models to guide or optimization and/or using more generic molecular representations (2D or 3D pharmacophore-based) to build the QSAR predictive models could be considered.

All in all, it is still uncertain to which extent AI-guided optimization can bring benefit to Lead Optimization in terms of reduction of number of compounds and number of iterations needed to identify a new molecular entity, and it would be of high interest to test such approach along several design-make-test cycles, starting early in the Lead Optimization phase, to assess the magnitude of the benefit brought by AI. Ideally, this should be conducted as a comparative "blind" study comparing the AI approach to a traditional approach to enable to draw strong conclusions.

Also, worthwhile mentioning, the selection process of the AI-designed molecules was not only based on pure data-driven ranking. Molecules were selected based on their scores on the predictive models, but also based on their synthetic accessibility and the expert input of medicinal chemists and computational chemists using their expertise as well as specific data visualization tools to remove poor quality compounds or potential false positives. This selection process associating human expertise and data visualization to rank and select AI-driven ideas was probably an important success factor in this project. Indeed, although not addressed in this paper, issues with synthetic accessibility, complexity, structural alerts issues, or sheer meaningless of certain AI propositions did occur in this project, although they remained minor in this context. These issues currently prevent a fully automated compound selection and rather advocate for a collaboration between chemists and AI, enabling to get the best of both worlds. Recently, notable progress has been made in the development of efficient methods for high throughput synthetic accessibility scoring^{35, 36, 37} which opens the perspective of an increased automation of the process.

6. CONCLUSION

Exploiting a sparse dataset of 881 molecules measured on 11 bioactivity assays, a DL-based AI *de novo* design algorithm was able to generate 150 virtual compounds with optimal *in silico* profiles against all desired characteristics of the project's TPP. Amongst those, 11 compounds were synthesized and measured on all 11 criteria of the TPP. The AI-designed molecules outperformed the ones designed by traditional medicinal chemistry approaches, achieving superior MPO scores. More importantly, three of those were found to meet the project's TPP, one of them strictly meeting all MPO objectives, the other two matching 10 objectives and being in the error margin of the assay for the last one. The AI algorithm came up with functional groups, which, although being rare or absent in the initial dataset, turned out to be highly beneficial for the MPO.

To our knowledge, this is the first report of a successful application of deep learning to *de novo* design for solving an MPO issue in an actual drug discovery project, moreover on a large number of objectives. This brings unequivocal evidence of the potential of this technology to bring substantial improvements to medicinal chemistry. The use of such approach in earlier stages of drug discovery (*i.e.* hit discovery, hit to lead and early LO) is under investigation. Improvement needs have been identified and are being addressed, notably regarding synthetic accessibility, compound complexity and domain of applicability of the predictive models.

7. CONFLICT OF INTEREST

Yann Gaston-Mathé, Nicolas Do-Huu and Quentin Perron are co-founders of Iktos (iktos.ai) and respectively CEO, CTO and CSO of the company. Hamza Tajmouati and Adam Skiredj are employees of Iktos.

Olivier Mirguet, Anne Rojas, Arnaud Gohier, Pierre Ducrot, Marie-Pierre Bourguignon, Patricia Sansilvestri-Morel and Françoise Gellibert are all employees of Servier.

8. REFERENCES

- (1) Bohacek, R.S., McMartin, C. and Guida, W.C. The art and practice of structure-based drug design: A molecular modeling perspective. *Med. Res. Rev.* **1996**, *16*, 3-50.
- (2) Lipinski, C. & Hopkins, A. Navigating chemical space for biology and medicine. *Nature*, **2004**, *432*, 855–861.
- (3) Nicolaou, C. A., Brown, N. Multi-Objective Optimization Methods in Drug Design. *Drug Discov. Today Technol.* **2013**, *10*, e427-e435.
- (4) Firth, N. C., Atrash, B., Brown, N. and Blagg, J. MOARF, an Integrated Workflow for Multiobjective Optimization: Implementation, Synthesis, and Biological Evaluation. *J. Chem. Inf. Model.* **2015**, *55*, 1169-1180.
- (5) Paul, S. M., Mytelka, D. S., Dunwiddie, C. T., Persinger, C. C., Munos, B. H., Lindborg, S. R. and Schacht, A. L. How to improve R&D productivity: the pharmaceutical industry's grand challenge. *Nat. Rev. Drug Discov.* **2010**, *9*, 203–214.
- (6) Talevi, A. Computer-Aided Drug Design: An Overview. **2018** In: Gore M., Jagtap U. (eds) Computational Drug Discovery and Design. Methods in Molecular Biology, vol 1762. Humana Press, New York, NY.
- (7) Kitchen, D., Decornez, H., Furr, J. et al. Docking and scoring in virtual screening for drug discovery: methods and applications. *Nat. Rev. Drug. Discov.* **2004**, *3*, 935–949.
- (8) Lavecchia, A. and Di Giovanni, C. Virtual Screening Strategies in Drug Discovery: A Critical Review. *Curr. Med. Chem.* **2013**, *20*, 2839-2860.
- (9) Neves, B. J., Braga, R. C., Melo-Filho, C. C., Moreira-Filho, J. T., Muratov, E. N., & Andrade, C. H. QSAR-Based Virtual Screening: Advances and Applications in Drug Discovery. *Front. Pharmacol.* **2018** *9*, 1275.
- (10) Nicolaou, C. A., Apostolakis, J. and Pattichis, C. S. De Novo Drug Design Using Multiobjective Evolutionary Graphs. *J. Chem. Inf. Model.* **2009**, *49*, 295–307.
- (11) Schneider, P., Walters, W. P., Plowright, A. T., Sieroka, N., Listgarten, J., Goodnow, R. A., Fisher, J., Jansen, J., Duca, J., Rush, T., Zentgraf, M., Hill, J. E., Krutoholow, E., Kohler, M., Blaney, J., Funatsu, K., Luebke, C. and Schneider, G. Rethinking drug design in the artificial intelligence era. *Nat. Rev. Drug. Discov.* **2020**, *19*, 353–364.
- (12) Weininger, D., SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *J. Chem. Inf. Comput. Sci.* **1988**, *28*, 31-36.
- (13) Segler, M.H.S., Kogej, T., Tyrchan, C., Waller, M.P.: Generating Focussed Molecule Libraries for Drug Discovery with Recurrent Neural Networks. *ACS Cent. Sci.* **2018**, *4*, 120-131.
- (14) Krenn, M., Häse, F., Nigam, A., Friederich, P. and Aspuru-Guzik A. Self-Referencing Embedded Strings (SELFIES): A 100% robust molecular string representation. *ArXiv*, **2019**, DOI: arXiv:1905.13741v2.
- (15) Nigam, A., Friederich, P., Krenn, M. and Aspuru-Guzik A. Augmenting Genetic Algorithms with Deep Neural Networks for Exploring the Chemical Space. *ArXiv*, **2019**, DOI: arXiv:1909.11655v4.

- (16) O'Boyle, N. and Dalke, A. DeepSMILES: An Adaptation of SMILES for Use in Machine-Learning of Chemical Structures. *ChemRxiv*, **2018**, DOI : 10.26434/chemrxiv.7097960.v1.
- (17) Zhou, Z., Kearnes, S., Li, L., Zare, R. N. and Riley, P. Optimization of Molecules via Deep Reinforcement Learning. *ArXiv*, **2019**, DOI: arXiv:1810.08678v3.
- (18) Popova, M., Shvets, M., Oliva, J. and Isayev, O. MolecularRNN: Generating realistic molecular graphs with optimized properties. *ArXiv*, **2019**, DOI: arXiv:1905.13372.
- (19) Jin, W., Barzilay, R. and Jaakkola, T. Junction Tree Variational Autoencoder for Molecular Graph Generation. *ArXiv*, **2018**, DOI: arXiv:1802.04364v4.
- (20) Gómez-Bombarelli, R., Wei, J. N., Duvenaud, D., Hernández-Lobato, J. M., Sánchez-Lengeling, B., Sheberla, D. and Aspuru-Guzik, A. Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. *ACS Cent. Sci.* **2018**, *4*, 268–276.
- (21) Guimaraes, G., Sanchez-Lengeling, B., Outeiral, C., Cunha Farias, P. and Aspuru-Guzik, A. Objective-Reinforced Generative Adversarial Networks (ORGAN) for Sequence Generation Models. *ArXiv*, **2017**, DOI: arXiv:1705.10843v3.
- (22) Olivecrona, M., Blaschke, T., Engkvist, O. and Chen, H. Molecular de-novo design through deep reinforcement learning. *J. Cheminform.* **2017**, *9*, 48.
- (23) Neil, D., Segler, M.H., Guasch, L., Ahmed, M., Plumbley, D., Sellwood, M., & Brown, N. Exploring Deep Recurrent Models with Reinforcement Learning for Molecule Design. *ICLR* **2018**.
- (24) Kusner, M. J., Paige, B. and Hernández-Lobato, J. M. Grammar Variational Autoencoder. *ArXiv*, **2017**, DOI: arXiv:1703.01925v1.
- (25) Winter, R., Montanari, F., Steffen, A., Briem, H., Noé, F., & Clevert, D.-A. Efficient multi-objective molecular optimization in a continuous latent space. *Chem. Sci.*, **2019**, *10*, 8016–8024.
- (26) Schneider, G. Automating drug discovery. *Nat. Rev. Drug Discov.* **2018**, *17*, 97.
- (27) Rogers, D., & Hahn, M. Extended-Connectivity Fingerprints. *J. Chem. Inf. Model.* **2010**, *50*, 742–754.
- (28) Bradley, A. P. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognit.* **1997**, *30*, 1145–1159.
- (29) Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation* **1997**, *9*, 1735–1780.
- (30) Williams, R. J. and Zipser, D. A Learning Algorithm for Continually Running Fully Recurrent Neural Networks. *Neural Computation*, **1989**, *1*, 270–280.
- (31) Bickerton, G. R., Paolini, G. V., Besnard, J., Muresan, S., Andrew, L. and Hopkins, A. L. Quantifying the chemical beauty of drugs. *Nat. Chem.* **2012**, *4*, 90–98.
- (32) Eckert, H. and Bajorath, J. Molecular similarity analysis in virtual screening: foundations, limitations and novel approaches. *Drug Discov. Today* **2007**, *12*, 225–233.
- (33) Leeson, P. D., & Young, R. J. Molecular Property Design: Does Everyone Get It? *ACS Med. Chem. Lett.* **2015**, *6*, 722–725.
- (34) Wold, S., Esbensen, K., & Geladi, P. Principal component analysis. *Chemometrics and Intelligent Laboratory Systems*, **1987**, *2*, 37–52.
- (35) Coley, C. W., Rogers, L., Green, W. H. and Jensen, K. F. SCScore: Synthetic Complexity Learned from a Reaction Corpus. *J. Chem. Inf. Model.* **2018**, *58*, 252–261.
- (36) Thakkar, Amol; Chadimova, Veronika; Bjerrum, Esben Jannik; Engkvist, Ola; Reymond, Jean-Louis (2020): Retrosynthetic Accessibility Score (RAscore) - Rapid Machine Learned Synthesizability Classification from AI Driven Retrosynthetic Planning. *ChemRxiv*. Preprint. <https://doi.org/10.26434/chemrxiv.13019993.v1>

(37) https://iktos.ai/wp-content/uploads/2020/09/Poster-Gen-Access-Synth-RSC-09_2020_-final.pdf