

# Beam search sampling for molecular design and intrinsic prioritization with machine intelligence

Michael Moret<sup>1#</sup>, Moritz Helmstädter<sup>2#</sup>, Francesca Grisoni<sup>1</sup>, Gisbert Schneider<sup>1\*</sup>, Daniel Merk<sup>2\*</sup>

<sup>1</sup> Department of Chemistry and Applied Biosciences, ETH Zurich, Vladimir-Prelog-Weg 4, 8093 Zurich, Switzerland

<sup>2</sup> Institute of Pharmaceutical Chemistry, Goethe University Frankfurt, Max-von-Laue-Str. 9, 60438 Frankfurt, Germany

# M.M. and M.H. contributed equally to this study.

\* gisbert@ethz.ch (G.S.), merk@pharmchem.uni-frankfurt.de (D.M.)

**Abstract.** Chemical language models enable de novo drug design without the requirement for explicit molecular construction rules. While such models have been applied to generate novel compounds with desired bioactivity, the actual prioritization and selection of the most promising computational designs remains challenging. In this work, we leveraged the probabilities learnt by chemical language models with the beam search algorithm as a model-intrinsic technique for automated molecule design and scoring. Prospective application of this method yielded three novel inverse agonists of retinoic acid receptor-related orphan receptors (RORs). Each design was synthesizable in three reaction steps and presented low-micromolar to nanomolar potency towards ROR $\gamma$ . This model-intrinsic sampling technique eliminates the strict need for external compound scoring functions, thereby further extending the applicability of generative artificial intelligence to data-driven drug discovery.

**Keywords.** de novo design, deep learning, drug discovery, neural network, nuclear receptor

## Introduction

Generative deep learning<sup>1,2</sup> can be applied to design pharmacologically active compounds de novo<sup>3–5</sup>. Deep learning-based molecular design algorithms can extract high-level molecular features from “raw” molecular representations<sup>6–9</sup>, such as molecular graphs and the Simplified Molecular Input Line Entry System (SMILES)<sup>10</sup>, potentially allowing them to access unexplored regions of the chemical space<sup>11</sup>. Previous studies demonstrated that “chemical language models” (CLMs)<sup>12,13</sup> that were trained on SMILES strings can generate novel molecules with experimentally validated bioactivity<sup>9,14</sup>. CLMs have shown the ability to learn focused chemical features from small collections of template molecules thanks to transfer learning which reuses previously learned knowledge on a new task<sup>14–16</sup>. A CLM is first pre-trained on a large amount of unlabeled data to capture the SMILES ‘syntax’ (*i.e.*, how alphanumeric characters should be assembled to generate strings that correspond to valid molecules) and the physicochemical properties of the pre-training dataset. Then, the pre-trained CLM is “fine-tuned” with a smaller set of task-specific data<sup>12,17,18</sup>. During this transfer learning process the CLM is biased towards the chemical space of interest, that is, molecules with desired biological and physicochemical properties. This “few-shot”<sup>19,20</sup> learning ability renders CLMs particularly useful for application to biological targets, for which only few ligands are known.

Previous prospective applications of CLMs for de novo molecule generation relied on “temperature sampling” of virtual molecule libraries. Temperature sampling generates molecules by weighted random sampling from the probability distributions learnt by the CLM during training. Such generated SMILES strings might not always be chemically meaningful (invalid strings) or they do not match the feature distribution of the training data. Additional methods are usually needed to select the most promising designs from the virtual molecule libraries, *e.g.*, based on the similarity to known bioactive molecules or activity predictions<sup>9,12,14</sup>. Here, we used the beam search algorithm as a model-intrinsic alternative to temperature sampling<sup>21,22</sup>. This method enables the automatic generation and prioritization of the designs, without strictly requiring additional compound scoring. Beam search scoring was successfully validated in a prospective application aiming to generate new retinoic acid related orphan receptor (ROR)<sup>23</sup> ligands from scratch.


RORs were chosen as molecular targets because these receptor proteins are an attractive but not extensively studied family of potential drug targets. They constitute a family of ligand-activated transcription factors that mainly act as monomers involved in the circadian control of energy homeostasis<sup>24,25</sup> and immune system regulation<sup>26,27</sup>, among other functions. RORs hold promising pharmacological potential for various indications, specifically for autoimmune diseases<sup>26,27</sup>. No ROR ligand has reached drug approval to date, partially owing to compound-related issues such as poor aqueous solubility, lack of selectivity, and clinical safety concerns<sup>26,28,29</sup>.

## Results & Discussion

### *Chemical language model and beam search sampling for de novo design*

We explored the beam search algorithm<sup>30</sup> to generate molecules from a CLM as a potential alternative to temperature sampling combined with an external ranking method. Given the probabilities learnt by a CLM, a vast number of SMILES strings could be sampled. As it is computationally not feasible to sample all outputs, a heuristic method such as beam search can be used to find the likely outputs. Here, the motivation to find the most likely SMILES strings was based on the hypothesis that the probability for generating a certain SMILES string correlates with the quality of the corresponding molecule with regard to the implicit design objective encoded in the fine-tuning set (e.g., desired bioactivity, physicochemical properties). During molecule generation by beam search sampling, the algorithm progressively adds characters (“tokens”) to a SMILES string while keeping track of the  $k$  most likely SMILES string(s). When adding a new token, the algorithm computes the probability of each possible growing SMILES string to define the  $k$  most likely SMILES string(s) for the next step (Fig 1c). This process is repeated until the SMILES sequence is completed (*i.e.*, the ‘end-of-sequence’ character is added) or a predefined maximal string length is reached. Thus, beam search can be used to approximately find the most probable molecules, as predicted by (i) the underlying model and (ii) the scoring function, which is computed as the product of the probabilities of each character (Fig. 1c). This scoring function (“beam search score”) allows to rank the de novo designs according to the probability of their SMILES tokens.

As a framework to probe beam search sampling, we employed a recently published CLM based on a recurrent neural network with long short-term memory (LSTM)<sup>31</sup> cells. The CLM was trained on 365,063 molecules from ChEMBL<sup>32</sup>, encoded as SMILES strings (Fig. 1a), to iteratively predict the next character of each SMILES string, given the preceding tokens (Fig. 1b). This pre-trained CLM was then biased towards the design objective, *i.e.*, the generation of new molecules with bioactivity on RORs, by transfer learning using collections of known ROR ligands (Supplementary Figure 1, Supplementary Table 1).



**Figure 1 | Molecule generation with a chemical language model (CLM) and beam search sampling.** (a) Kekulé structure of an example molecule and corresponding SMILES string. (b) CLM training. The CLM learns to predict the probability of each character of a SMILES string based on the sequence of the respective previous characters. (c) Beam search decoding of width two ( $k = 2$ ): The design algorithm keeps track of the two most likely SMILES strings (highlighted in color). In this example, the SMILES string generation proceeds from left to right.


### *Application of beam search sampling to designing inverse ROR $\gamma$ agonists*

For prospective evaluation, we applied the beam search to the design of natural-product inspired ROR $\gamma$  ligands. Learning from natural products as a traditional source of inspiration for drug discovery<sup>33,34</sup> may hold several advantages over learning from purely synthetic molecules, because of the overall higher structural diversity, greater three-dimensionality, and often superior selectivity of bioactive natural products<sup>35,36</sup>. We aimed to obtain de novo designs possessing three properties: (i) natural product inspired chemical structure, (ii) synthesizability, and (iii) bioactivity on ROR $\gamma$ . Aiming to fulfil all three objectives during transfer learning, the CLM that was previously pre-trained on bioactive molecules from ChEMBL<sup>15</sup> was fine-tuned on one synthetic and four natural product ROR $\gamma$  modulators (Suppl. Fig. 1). Beam search sampling was started after the fifth epoch of fine-tuning, to ensure that the CLM had sufficiently captured the molecular features of the small fine-tuning set.

All valid SMILES strings generated between epochs 5 and 16 (last fine-tuning epoch) were ranked by beam search scoring. The top five designs according to the beam search score (Fig. 2a) were deemed synthetically inaccessible by medicinal chemists.

This was further highlighted by a machine learning algorithm for retrosynthetic analysis (IBM RXN)<sup>37</sup> which did not find a synthetic route for any of these molecules. Apparently, the CLM failed to meet the fundamental design criterion of synthesizability, suggesting that a different fine-tuning strategy might be required.

These findings show a first important benefit of the beam search sampling technique. Beam search sampling reveals the most likely designs from a CLM which can be assessed regarding their compliance with the design objectives to evaluate model performance and the success of fine-tuning.

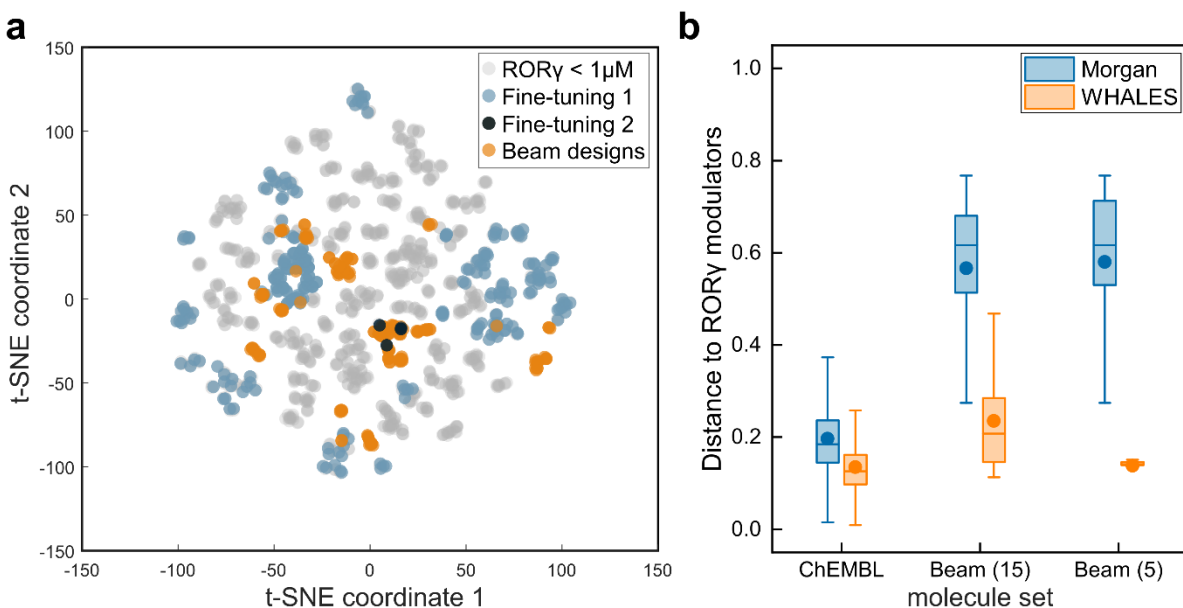


**Figure 2 | Top ranking designs obtained by beam search sampling.** (a) Single fine-tuning, (b) double fine-tuning. Ranks are based on the beam search score of the molecular designs. For the top ranking molecules from the double fine-tuning experiment, the similarity values refer to the Jaccard-Tanimoto similarity computed on Morgan fingerprints (length=1024, 2-bond radius) to the closest known active molecule annotated in ChEMBL with an IC<sub>50</sub> value for ROR $\gamma$ .

Aiming to improve upon these results, we performed a second experiment, in which we applied a stepwise fine-tuning strategy. First, the pre-trained model was fine-tuned for 20 epochs on 255 synthetic ROR $\gamma$  ligands from the US patent subset of the Protein Data Bank<sup>38</sup> (255 molecules, Suppl. Table 1) to capture both bioactivity and synthesizability. Then, the model was fine-tuned with four natural product ROR $\gamma$  modulators (Suppl. Fig. 1) for 16 epochs, aiming to bias the model towards natural-product-likeness. With this second approach, the top 5 sampled molecules (Figure 2b) were synthetically accessible according to IBM RXN<sup>37</sup>, which could propose a synthetic route for each of them. Importantly, the computer-generated molecules possess certain natural product characteristics (Fig. 3, Suppl. Table 2), as indicated by a high fraction of sp<sup>3</sup>-hybridized carbon atoms (Fsp<sup>3</sup>). The top five designs have Fsp<sup>3</sup> values ranging from 50% to 75%. These values are comparable to the values computed for the MEGx natural product library (Analyticon Discovery GmbH, rel. 09-01-2018), and exceed the average Fsp<sup>3</sup> content of the ChEMBL molecules used for pre-training (51 $\pm$ 30% and 33 $\pm$ 20%, respectively). These

data suggested that the two-step fine-tuning procedure complied with the design objectives. The two-step approach was chosen for prospective application.

We then compared the beam search designs obtained with the chosen computational strategy to known ROR $\gamma$  modulators and to the fine-tuning molecules (Fig. 3). Despite being a “greedy” algorithm, the beam search sampling still allowed to explore the chemical space beyond the regions that are populated by the fine-tuning compounds (Figure 3a). Compared to the inverse ROR $\gamma$  agonists annotated in ChEMBL (IC<sub>50</sub> <1  $\mu$ M, Figure 3b), the beam search designs are structurally more diverse in terms of substructure fragments encoded by Morgan fingerprints<sup>39</sup>. The designs possess a certain degree of similarity to the known active molecules in terms of their three-dimensional shape and partial charge distribution (WHALES descriptors<sup>40</sup>). Apparently, the CLM, on top of learning the SMILES “syntax”, also learned certain “semantic” ligand features that are relevant for binding to macromolecules, such as their molecular shape and partial charge patterns.




**Figure 3 | Characteristics of designs from the balanced CLM.** (a) Stochastic neighbor embedding (t-SNE)<sup>41</sup> projection of the compounds sets based on Morgan fragment fingerprints (length = 1024, 2-bond radius, Jaccard-Tanimoto similarity). The location of the two-fine tuning sets, the ROR $\gamma$  modulators annotated in ChEMBL (IC<sub>50</sub> <1  $\mu$ M, 1091 compounds), and the beam search designs are shown. (b) Comparison of the sampled molecular designs with known ROR $\gamma$  modulators (IC<sub>50</sub> <1  $\mu$ M) in terms of Morgan fragment fingerprints (“Morgan”) and three-dimensional shape and electrostatics descriptors (WHALES). The pairwise distance distribution among known ChEMBL modulators is shown as a reference. “Beam (15)” and “Beam (5)” indicate the top 15 and top 5 designs, respectively. Boxplots indicate 25<sup>th</sup>, 50<sup>th</sup>, and 75<sup>th</sup> percentiles (lines), mean values (circle), and outlier boundaries (whiskers, 1.5  $\times$  interquartile range).

### *Prospective experimental validation*

Three beam search designs were synthesized and characterized in vitro. We selected them based on their beam search score. From the five most likely designs (Figure 2b), we selected molecules **1** and **2**, which were ranked first and third. Compound **2** showed the highest Tanimoto similarity (Morgan fingerprints) to a known ROR $\gamma$  modulator (Figure 2b). The scaffolds of both compounds were also prominent among the beam search designs not included in the top 5, suggesting a structural preference. The scaffold of **1** also appeared in the design ranked 6<sup>th</sup>. Molecules ranked 10<sup>th</sup> and 13<sup>th</sup> resembled compound **2**. Hence, we additionally chose compound **3** of this abundant chemotype from rank 13 for prospective validation. Compounds **1-3** were synthesized according to Scheme 1.

For preparation of **1**, (4-chlorophenyl)-piperazine (**4**) was treated with 4-bromobutyl acetate (**5**) to obtain the ester protected intermediate **6** which after alkaline ester hydrolysis to **7** was suitable for Mitsunobu reaction with 8-azaspiro[4.5]decane-7,9-dione (**8**) to obtain the top-ranking computational design **1**. Preparation of **2** started from 4-bromo-2-fluorobenzaldehyde (**9**) which was reacted with amine **10** to obtain **11** by reductive amination followed by sulfonamide coupling with **12** to yield **13**. Eventually, the 4-trifluoromethylpiperidine substituent was introduced to **13** under Buchwald-Hartwig conditions with **14** yielding compound **2**. The structurally related design **3** was not accessible via this route and was prepared in a different reaction sequence starting from 4-fluorobenzaldehyde (**15**), which was reacted with 4-trifluoromethylpiperidine (**14**) to **16** in a nucleophilic aromatic substitution. Reductive amination with cyclobutaneamine (**10**) to **17**, followed by sulfonamide formation with phenylmethanesulfonyl chloride (**12**), afforded the computationally designed compound **3**.

In vitro characterization of compounds **1**, **2**, and **3** in Gal4-ROR hybrid reporter gene assays confirmed inverse ROR $\gamma$  agonism with micromolar to sub-micromolar IC<sub>50</sub> values (Table 1). The top-ranking compound **1** counteracted ROR $\gamma$  activity with an IC<sub>50</sub> value of 4.6  $\mu$ M. It was additionally active on ROR $\alpha$  and ROR $\beta$ , but precise IC<sub>50</sub> values could not be determined due to cytotoxicity. Compounds **2** and **3** counteracted ROR $\gamma$  activity with IC<sub>50</sub> values of 0.37  $\mu$ M (**2**) and 0.68  $\mu$ M (**3**), respectively. In addition to being inverse ROR $\gamma$  agonists, all three synthesized designs revealed pronounced preference for the ROR $\gamma$  subtype, with compounds **2** and **3** possessing more than tenfold higher potency on ROR $\gamma$  compared to the related ROR $\alpha$  and ROR $\beta$  isoforms. These results show that the CLM with beam search sampling conserved the bioactivity of the training molecules in the computational designs.



**Scheme 1 |** Synthesis of the CLM designs **1**, **2** and **3**. Reagents & Conditions: (a) DMF, 4-DMAP, 60°C, 16 h, 48%; (b) KOH, H<sub>2</sub>O/THF/MeOH,  $\mu$ w, 100°C, 30 min, 98%; (c) DIAD, PPh<sub>3</sub>, THF, 0°C to rt, 16 h, 42%; (d) NaB(OAc)<sub>3</sub>H, HOAc, DCE, r.t., 50 h, 73%; (e) 4-DMAP, pyridine, CH<sub>2</sub>Cl<sub>2</sub>, reflux, 16 h, 37%; (f) Pd<sub>2</sub>(dba)<sub>3</sub>, xantphos, Cs<sub>2</sub>CO<sub>3</sub>, 1,4-dioxane, reflux, 16 h, 18%; (g) K<sub>2</sub>CO<sub>3</sub>, DMSO, reflux, 48 h, 82%.

**Table 1 |** Activity of de novo designs **1**, **2** and **3** on RORs in Gal4 hybrid reporter gene assays. Data are the mean $\pm$ S.E.M.,  $n \geq 4$ .

ID	structure	IC <sub>50</sub> (ROR $\alpha$ )	IC <sub>50</sub> (ROR $\beta$ )	IC <sub>50</sub> (ROR $\gamma$ )
<b>1</b>		>10 $\mu$ M	>10 $\mu$ M	4.6 $\pm$ 0.5 $\mu$ M
<b>2</b>		23 $\pm$ 3 $\mu$ M	22 $\pm$ 1 $\mu$ M	0.37 $\pm$ 0.05 $\mu$ M
<b>3</b>		10 $\pm$ 1 $\mu$ M	7.6 $\pm$ 0.5 $\mu$ M	0.68 $\pm$ 0.07 $\mu$ M



## Conclusion

In this work, beam search sampling from CLMs was applied to generating new molecules with desired bioactivity on the ligand-activated transcription factor ROR $\gamma$ . The algorithm successfully generated and scored new molecules. Prospective experimental validation yielded three novel, potent inverse agonists of the nuclear receptor. The selected compounds have various degrees of similarity to known ROR $\gamma$  modulators (ranging from 0.28 to 0.71, as captured by Jaccard-Tanimoto similarity on Morgan fingerprints). Apparently the beam search approach coupled with a CLM conserves structural features necessary for the desired bioactivity but still generates innovative and structurally diverse compounds. This result corroborates beam search sampling as a technique for de novo design of bioactive molecules by a CLM. The computational and experimental results suggest two advantages of the beam search algorithm that could further the applicability of AI for de novo design. Firstly, by searching for the most likely molecules a CLM can generate with greedy decoding, the beam search algorithm can serve as an internal measure to probe the suitability of a CLM for the task under analysis. Analysis of the beam search designs can serve to analyze compliance of the CLM designs with the design objectives and hence to evaluate the success of fine-tuning. This is in contrast with temperature sampling, which might lead chemists to consider designs that are not likely according to the model. Beam search may thus be considered an alternative to temperature sampling. It should be noted, however, that the number of designs that can be sampled by beam search is limited compared to temperature sampling, which can virtually generate an infinite number of chemical structures. The two techniques complement each other as alternatives, each with characteristic advantages. The desired application should guide the choice of either strategy. Secondly, beam search sampling could potentially overcome the need for external compound prioritization. If corroborated in future prospective studies, this feature may help to further the applicability of CLMs for de novo molecular design in medicinal chemistry.

## Methods

*Data processing.* Molecules were encoded as canonical SMILES strings<sup>10</sup> using the RDKit package (v.2018.03, [www.rdkit.org](http://www.rdkit.org)) and only SMILES strings with a length of up to 140 characters were retained. SMILES strings were standardized in Python (v3.6.5, [www.python.org](http://www.python.org)) by removing stereochemical information, salts and duplicates.

*Datasets.* We used the processed data ([https://github.com/ETHmodlab/virtual\\_libraries](https://github.com/ETHmodlab/virtual_libraries))<sup>15</sup> we recently published for both ChEMBL (used for pretraining the CLM) and for the representative natural products library, MEGx (released 01 September 2018, Analyticon Discovery GmbH). In total, the processed version of ChEMBL contains 365,063 molecules, and the processed version of MEGx 2,931 molecules. The 255 modulators of

RORy used for double fine-tuning were taken from the US patent subset of the protein data bank (rcsb.org)<sup>38</sup>, and the four natural products, along with the synthetic compound were taken from the study of Solt and Burris<sup>42</sup>. The most similar known RORy modulators with reported IC<sub>50</sub> were extracted from ChEMBL (organism: *Homo sapiens*).

*Chemical language model.* We used our recently published framework to implement the CLM ([https://github.com/ETHmodlab/virtual\\_libraries](https://github.com/ETHmodlab/virtual_libraries))<sup>15</sup>, which is based on a long short-term memory (LSTM)<sup>31</sup>. The model implementations differed based on the transfer learning strategy used: (i) For single fine-tuning, the model was composed of four layers that have a total of 5,820,515 parameters (layer 1: BatchNormalization, layer 2: LSTM with 1,024 units, layer 3: LSTM with 256 units and layer 4: BatchNormalization). We pretrained the CLM for 10 epochs with a learning rate of  $10^{-3}$  and performed the fine-tuning for 16 epochs with a learning rate of  $10^{-4}$  by keeping layer 2 frozen. (ii) For double fine-tuning, the model was composed of five layers that have a total of 8,444,003 parameters (layer 1: BatchNormalization, layer 2: LSTM with 1,024 units, layer 3: LSTM with 512 units and layer 4: LSTM with 256 units, layer 5: BatchNormalization). We pretrained the CLM for 10 epochs with a learning rate of  $10^{-3}$  and performed two rounds of fine-tuning. The first round was applied for 20 epochs with a learning rate of  $10^{-4}$  by keeping layer 2 frozen. The second round was applied for 20 epochs with a learning rate of  $10^{-4}$  by keeping layer 2 and 3 frozen. The CLMs for both strategies were trained on SMILES strings encoded as one-hot vectors. We used the categorical cross-entropy loss and the Adam optimizer<sup>43</sup>. We applied a 10-fold data augmentation to all molecules.

*Beam search ranking.* We applied the following procedure for the beam search algorithm:

---

**Algorithm 1:** Beam search to sample SMILES from a CLM

---

**Result:** List of generated SMILES

Beam width =  $k$ ;

Maximum SMILES length =  $l$ ;

Initialize list of potential SMILES with starting token:  $potential_{smi} = ['G']$ ;

**while** all SMILES in  $potential_{smi}$  are not complete or number of loop <  $l$  **do**

**for each** SMILES in  $potential_{smi}$  **do**

**for each character in the vocabulary** **do**

            Get probabilities of adding the character to the SMILES by the CLM;

            Create new SMILES by adding the character to the SMILES;

            Compute overall score of the SMILES by multiplying probabilities of each of its character;

**end**

**end**

    Rank list of new SMILES by most likely to less likely according to overall score;

    Prune list of potential SMILES to the  $k$  most likely candidates;

    Replace list of  $potential_{smi}$  by the  $k$  most likely candidates;

**end**

---

We used a beam width ( $k$ ) of 50 and defined the maximum SMILES strings length ( $l$ ) as 140 tokens. We considered in the ranking molecules sampled from epoch 5 to epoch 16

of fine-tuning; the first epochs were not used to ensure the model was sufficiently biased toward the fine-tuning set.

*Stochastic Neighbor Embedding.* The t-SNE projection was performed with the “tsne” function of MATLAB R2020a<sup>44</sup> on Morgan Fingerprints (‘Distance’ = ‘jaccard’). The perplexity value was optimized from 5 to 50 with a step equal to 10 to minimize the compression error. A perplexity equal to 5 was chosen, corresponding to the minimum compression error (0.328).

*Molecular descriptors.* Molecular geometry and the corresponding WHALES descriptors were computed with the code freely available at [https://github.com/grisoniFr/scaffold\\_hopping\\_whales](https://github.com/grisoniFr/scaffold_hopping_whales) (v1), with default settings, as previously described<sup>45</sup>. Morgan fingerprints<sup>39</sup> (length=1024, 2-bond radius) were computed using RDKit (v.2018.03) in Python (v3.6.5).

*IBM RXN.* We used the web interface with default parameters (<https://rxn.res.ibm.com/>).

Synthetic and in vitro pharmacological procedures are described in the Supplementary Information.

### Competing interest

G.S. declares a potential financial conflict of interest as a founder of inSili.com GmbH, Zurich, and in his role as consultant to the pharmaceutical industry.

### Acknowledgements

This research was supported by the Swiss National Science Foundation (grant no. 205321\_182176 to G.S.), the RETHINK initiative at ETH Zurich and the Novartis Forschungsstiftung (FreeNovation grant “AI in Drug Discovery” to G.S.).

### References

- (1) Schmidhuber, J. Deep Learning in Neural Networks: An Overview. *Neural Networks* **2015**, 61, 85–117. <https://doi.org/10.1016/j.neunet.2014.09.003>.
- (2) Lecun, Y.; Bengio, Y.; Hinton, G. Deep Learning. *Nature* **2015**, 521 (7553), 436–444. <https://doi.org/10.1038/nature14539>.
- (3) Chen, H.; Engkvist, O.; Wang, Y.; Olivecrona, M.; Blaschke, T. The Rise of Deep Learning in Drug Discovery. *Drug Discov. Today* **2018**, 23 (6), 1241–1250. <https://doi.org/10.1016/j.drudis.2018.01.039>.
- (4) Walters, W. P.; Barzilay, R. Applications of Deep Learning in Molecule Generation and Molecular Property Prediction. *Acc. Chem. Res.* **2020**, acs.accounts.0c00699. <https://doi.org/10.1021/acs.accounts.0c00699>.
- (5) Sanchez-Lengeling, B.; Aspuru-Guzik, A. Inverse Molecular Design Using Machine Learning: Generative Models for Matter Engineering. *Science* **2018**, 361

- (6400), 360–365. <https://doi.org/10.1126/science.aat2663>.
- (6) Gómez-Bombarelli, R.; Wei, J. N.; Duvenaud, D.; Hernández-Lobato, J. M.; Sánchez-Lengeling, B.; Sheberla, D.; Aguilera-Iparraguirre, J.; Hirzel, T. D.; Adams, R. P.; Aspuru-Guzik, A. Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. *ACS Cent. Sci.* **2018**, *4* (2), 268–276. <https://doi.org/10.1021/acscentsci.7b00572>.
  - (7) De Cao, N.; Kipf, T. MolGAN: An Implicit Generative Model for Small Molecular Graphs. *arXiv* **2018**.
  - (8) 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* (1–2), 1700111. <https://doi.org/10.1002/minf.201700111>.
  - (9) Merk, D.; Friedrich, L.; Grisoni, F.; Schneider, G. De Novo Design of Bioactive Small Molecules by Artificial Intelligence. *Mol. Inform.* **2018**, *37* (1), 1700153. <https://doi.org/10.1002/minf.201700153>.
  - (10) Weininger, D. SMILES, a Chemical Language and Information System: 1: Introduction to Methodology and Encoding Rules. *J. Chem. Inf. Comput. Sci.* **1988**, *28* (1), 31–36. <https://doi.org/10.1021/ci00057a005>.
  - (11) Skinnider, M. A.; Stacey, R. G.; Wishart, D. S.; Foster, L. J. Deep Generative Models Enable Navigation in Sparsely Populated Chemical Space. **2021**. <https://doi.org/10.26434/CHEMRXIV.13638347.V1>.
  - (12) Segler, M. H. S.; Kogej, T.; Tyrchan, C.; Waller, M. P. Generating Focused Molecule Libraries for Drug Discovery with Recurrent Neural Networks. *ACS Cent. Sci.* **2018**, *4* (1), 120–131. <https://doi.org/10.1021/acscentsci.7b00512>.
  - (13) Yuan, W.; Jiang, D.; Nambiar, D. K.; Liew, L. P.; Hay, M. P.; Bloomstein, J.; Lu, P.; Turner, B.; Le, Q. T.; Tibshirani, R.; Khatri, P.; Moloney, M. G.; Koong, A. C. Chemical Space Mimicry for Drug Discovery. *J. Chem. Inf. Model.* **2017**, *57* (4), 875–882. <https://doi.org/10.1021/acs.jcim.6b00754>.
  - (14) Merk, D.; Grisoni, F.; Friedrich, L.; Schneider, G. Tuning Artificial Intelligence on the de Novo Design of Natural-Product-Inspired Retinoid X Receptor Modulators. *Commun. Chem.* **2018**, *1* (1), 68. <https://doi.org/10.1038/s42004-018-0068-1>.
  - (15) Moret, M.; Friedrich, L.; Grisoni, F.; Merk, D.; Schneider, G. Generative Molecular Design in Low Data Regimes. *Nat. Mach. Intell.* **2020**, *2* (3), 171–180. <https://doi.org/10.1038/s42256-020-0160-y>.
  - (16) Awale, M.; Sirockin, F.; Stiefl, N.; Reymond, J. L. Drug Analogs from Fragment-Based Long Short-Term Memory Generative Neural Networks. *J. Chem. Inf. Model.* **2019**, *59* (4), 1347–1356. <https://doi.org/10.1021/acs.jcim.8b00902>.
  - (17) Yosinski, J.; Clune, J.; Bengio, Y.; Lipson, H. How Transferable Are Features in Deep Neural Networks? *Adv. Neural Inf. Process. Syst.* **2014**, *27*, 3320–3328.
  - (18) Peters, M.; Ruder, S.; Smith, N. A. To Tune or Not to Tune? Adapting Pretrained Representations to Diverse Tasks. *arXiv* **2019**. <https://doi.org/10.18653/v1/w19-4302>.
  - (19) Altae-Tran, H.; Ramsundar, B.; Pappu, A. S.; Pande, V. Low Data Drug Discovery with One-Shot Learning. *ACS Cent. Sci.* **2017**, *3* (4), 283–293. <https://doi.org/10.1021/acscentsci.6b00367>.
  - (20) Wang, Y.; Yao, Q.; Kwok, J.; Ni, L. M. Generalizing from a Few Examples: A Survey on Few-Shot Learning. *arXiv* **2019**.
  - (21) Schwaller, P.; Laino, T.; Gaudin, T.; Bolgar, P.; Hunter, C. A.; Bekas, C.; Lee, A.

- A. Molecular Transformer: A Model for Uncertainty-Calibrated Chemical Reaction Prediction. *ACS Cent. Sci.* **2019**, 5 (9), 1572–1583. <https://doi.org/10.1021/acscentsci.9b00576>.
- (22) Grechishnikova, D. Transformer Neural Network for Protein-Specific de Novo Drug Generation as a Machine Translation Problem. *Sci. Rep.* **2021**, 11 (1), 321. <https://doi.org/10.1038/s41598-020-79682-4>.
- (23) Benoit, G.; Cooney, A.; Giguere, V.; Ingraham, H.; Lazar, M.; Muscat, G.; Perlmann, T.; Renaud, J. P.; Schwabe, J.; Sladek, F.; Tsai, M. J.; Laudet, V. International Union of Pharmacology. LXVI. Orphan Nuclear Receptors. *Pharmacol. Rev.* **2006**, 58 (4), 798–836. <https://doi.org/10.1124/pr.58.4.10>.
- (24) Marciano, D. P.; Chang, M. R.; Corzo, C. A.; Goswami, D.; Lam, V. Q.; Pascal, B. D.; Griffin, P. R. The Therapeutic Potential of Nuclear Receptor Modulators for Treatment of Metabolic Disorders: PPAR $\gamma$ , RORs, and Rev-Erbs. *Cell Metab.* **2014**, 19 (2), 193–208. <https://doi.org/10.1016/j.cmet.2013.12.009>.
- (25) Hoon Kim, Y.; Lazar, M. A. Transcriptional Control of Circadian Rhythms and Metabolism: A Matter of Time and Space. *Endocr. Rev.* **2020**, 41 (5). <https://doi.org/10.1210/endrev/bnaa014>.
- (26) Pandya, V. B.; Kumar, S.; Sachchidanand; Sharma, R.; Desai, R. C. Combating Autoimmune Diseases with Retinoic Acid Receptor-Related Orphan Receptor- $\gamma$  (ROR $\gamma$  or ROR $\gamma$ c) Inhibitors: Hits and Misses. *J. Med. Chem.* **2018**, 61 (24), 10976–10995. <https://doi.org/10.1021/acs.jmedchem.8b00588>.
- (27) Solt, L. A.; Burris, T. P. Action of RORs and Their Ligands in (Patho)Physiology. *Trends Endocrinol. Metab.* **2012**, 23 (12), 619–627. <https://doi.org/10.1016/j.tem.2012.05.012>.
- (28) Asimus, S.; Palmér, R.; Albayaty, M.; Forsman, H.; Lundin, C.; Olsson, M.; Pehrson, R.; Mo, J.; Russell, M.; Carlert, S.; Close, D.; Keeling, D. Pharmacokinetics, Pharmacodynamics and Safety of the Inverse Retinoic Acid-Related Orphan Receptor  $\gamma$  Agonist AZD0284. *Br. J. Clin. Pharmacol.* **2020**, 86 (7), 1398–1405. <https://doi.org/10.1111/bcp.14253>.
- (29) Kojetin, D. J.; Burris, T. P. REV-ERB and ROR Nuclear Receptors as Drug Targets. *Nat. Rev. Drug Discov.* **2014**, 13 (3), 197–216. <https://doi.org/10.1038/nrd4100>.
- (30) Lowerre, B. T. The Harpy Speech Recognition System. *PhD Thesis Carnegie Mellon Univ. Pittsburgh.* **1976**.
- (31) Hochreiter, S.; Schmidhuber, J. Long Short-Term Memory. *Neural Comput.* **1997**, 9 (8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>.
- (32) 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. *Nucleic Acids Res.* **2014**, 42 (Database issue), D1083–D1090. <https://doi.org/10.1093/nar/gkt1031>.
- (33) Newman, D. J.; Cragg, G. M. Natural Products as Sources of New Drugs over the Nearly Four Decades from 01/1981 to 09/2019. *J. Nat. Prod.* **2020**, 83 (3), 770–803. <https://doi.org/10.1021/acs.jnatprod.9b01285>.
- (34) Rodrigues, T.; Reker, D.; Schneider, P.; Schneider, G. Counting on Natural Products for Drug Design. *Nat. Chem.* **2016**, 8 (6), 531–541. <https://doi.org/10.1038/nchem.2479>.

- (35) Ertl, P.; Schuffenhauer, A. Cheminformatics Analysis of Natural Products: Lessons from Nature Inspiring the Design of New Drugs. In *Progress in Drug Research*; Birkhäuser Basel, 2008; Vol. 66, pp 217–235. [https://doi.org/10.1007/978-3-7643-8595-8\\_4](https://doi.org/10.1007/978-3-7643-8595-8_4).
- (36) Ertl, P.; Roggo, S.; Schuffenhauer, A. Natural Product-Likeness Score and Its Application for Prioritization of Compound Libraries. *J. Chem. Inf. Model.* **2008**, *48* (1), 68–74. <https://doi.org/10.1021/ci700286x>.
- (37) Schwaller, P.; Petraglia, R.; Zullo, V.; Nair, V. H.; Haeuselmann, R. A.; Pisoni, R.; Bekas, C.; Iuliano, A.; Laino, T. Predicting Retrosynthetic Pathways Using Transformer-Based Models and a Hyper-Graph Exploration Strategy. *Chem. Sci.* **2020**, *11* (12), 3316–3325. <https://doi.org/10.1039/c9sc05704h>.
- (38) Berman, H. M.; Battistuz, T.; Bhat, T. N.; Bluhm, W. F.; Bourne, P. E.; Burkhardt, K.; Feng, Z.; Gilliland, G. L.; Iype, L.; Jain, S.; Fagan, P.; Marvin, J.; Padilla, D.; Ravichandran, V.; Schneider, B.; Thanki, N.; Weissig, H.; Westbrook, J. D.; Zardecki, C. The Protein Data Bank. *Acta Crystallogr. Sect. D Biol. Crystallogr.* **2002**, *58* (6 I), 899–907. <https://doi.org/10.1107/S0907444902003451>.
- (39) Morgan, H. L. The Generation of a Unique Machine Description for Chemical Structures-A Technique Developed at Chemical Abstracts Service. *J. Chem. Doc.* **1965**, *5* (2), 107–113. <https://doi.org/10.1021/c160017a018>.
- (40) Grisoni, F.; Merk, D.; Consonni, V.; Hiss, J. A.; Tagliabue, S. G.; Todeschini, R.; Schneider, G. Scaffold Hopping from Natural Products to Synthetic Mimetics by Holistic Molecular Similarity. *Commun. Chem.* **2018**, *1* (1), 44. <https://doi.org/10.1038/s42004-018-0043-x>.
- (41) van der Maaten, L.; Hinton, G. Visualizing Data Using T-SNE. *J. Mach. Learn. Res.* **2008**, *9*, 2579–2605.
- (42) Solt, L. A.; Burris, T. P. Action of RORs and Their Ligands in (Patho)Physiology. *Trends Endocrinol. Metab.* **2012**, *23* (12), 619–627. <https://doi.org/10.1016/j.tem.2012.05.012>.
- (43) Kingma, D. P.; Ba, J. L. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*; International Conference on Learning Representations, ICLR, 2015.
- (44) The MathWorks Inc., MATLAB 2017a, Natick, Massachusetts, U. S. The MathWorks Inc., MATLAB 2017b, Natick, Massachusetts, United States. Natick, Massachusetts 2017.
- (45) Grisoni, F.; Schneider, G. Molecular Scaffold Hopping via Holistic Molecular Representation. *Methods Mol. Biol.* **2020**, *in press*.