Drug Analogs from Fragment Based Long Short-Term Memory Generative Neural Networks

Mahendra Awale,[a)] Finton Sirockin,[b)] Nikolaus Stiefl[b)] and Jean-Louis Reymond[a)]*

[a)] *Department of Chemistry and Biochemistry, University of Bern, Freiestrasse 3, 3012 Bern, Switzerland,* e-mail: jean-louis.reymond@dcb.unibe.ch

[b)] *Novartis Institutes for Biomedical Research, CH-4002 Basel, Switzerland*

Abstract

Several recent reports have shown that long short-term memory generative neural networks (LSTM) of the type used for grammar learning efficiently learn to write SMILES of drug-like compounds when trained with SMILES from a database of bioactive compounds such as ChEMBL and can later produce focused sets upon transfer learning with compounds of specific bioactivity profiles. Here we trained an LSTM using molecules taken either from ChEMBL, DrugBank, commercially available fragments, or from FDB-17 (a database of fragments up to 17 atoms) and performed transfer learning to a single known drug to obtain new analogs of this drug. We found that this approach readily generates hundreds of relevant and diverse new drug analogs and works best with training sets of around 40,000 compounds as simple as commercial fragments. These data suggest that fragment-based LSTM offer a promising method for new molecule generation.

Introduction

New drug-like small molecules are constantly needed to feed the drug discovery pipeline.[1, 2] While the number of possible molecules is extremely large,[3, 4] one can narrow the search for such new small molecules by exploiting the accumulated knowledge on drug-target interactions.[5] In one such approach it was recently discovered that long short-term memory generative neural networks (LSTM) of the type used for grammar learning,[6, 7] trained with SMILES (Simplified Molecular Input Line Entry System)[8] representing organic compounds from ChEMBL,[9] a large annotated database of bioactive molecules, can generate new drug-like molecular structures, which can even be tailored to specific targets upon transfer learning with focused subsets of bioactive compounds.[10-14] The molecules generated by LSTM retain structural features from the parent molecules, which focuses the generation process on analogs with a higher probability of shared bioactivity, and provides an advantage in terms of synthesis planning because such close analogs may be easier to synthesize using routes known for the parent molecules.

Here we performed transfer learning with a single drug molecule to generate new analogs of this drug, an implementation of LSTM towards analog generation that is simpler than previously reported implementations towards this goal (Figure 1).[13] We studied the influence of the primary training set of molecules on the outcome of LSTM for 10 drugs covering a broad range of size and complexity from small molecule drugs to macrocyclic natural products (Figure 2). We were specifically interested to compare the effect of training with bioactive molecules such as those from ChEMBL[9] or DrugBank[15] with small fragment-sized compounds. We selected fragments either from commercial catalogs or from FDB17,[16] a database of theoretically possible fragments covering the entire chemical space up to 17 atoms.[17] Our data shows that LSTM training with fragment-sized molecules leads to new analogs as efficiently as if training is done with drug type molecules from ChEMBL or DrugBank. Excellent results are obtained by LSTM training with a relatively small set of approximately 40,000 molecules, which allows to cover a relevant portion of chemical diversity

at the scale of fragments. Our data suggest that fragment-based LSTM offers a promising method for new molecule generation.

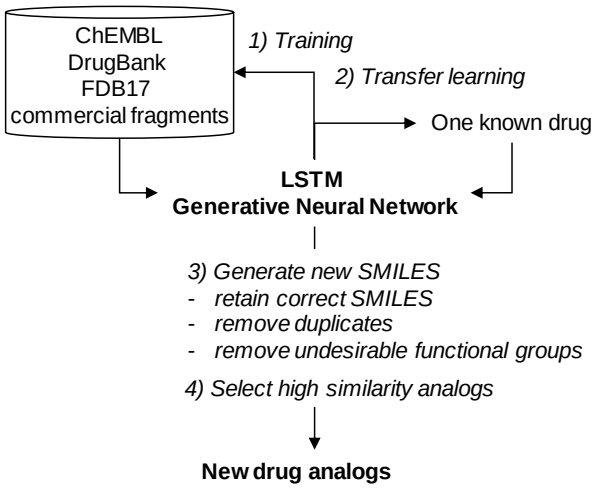

Figure 1. Principle of LSTM neural networks for generating drug analogs.

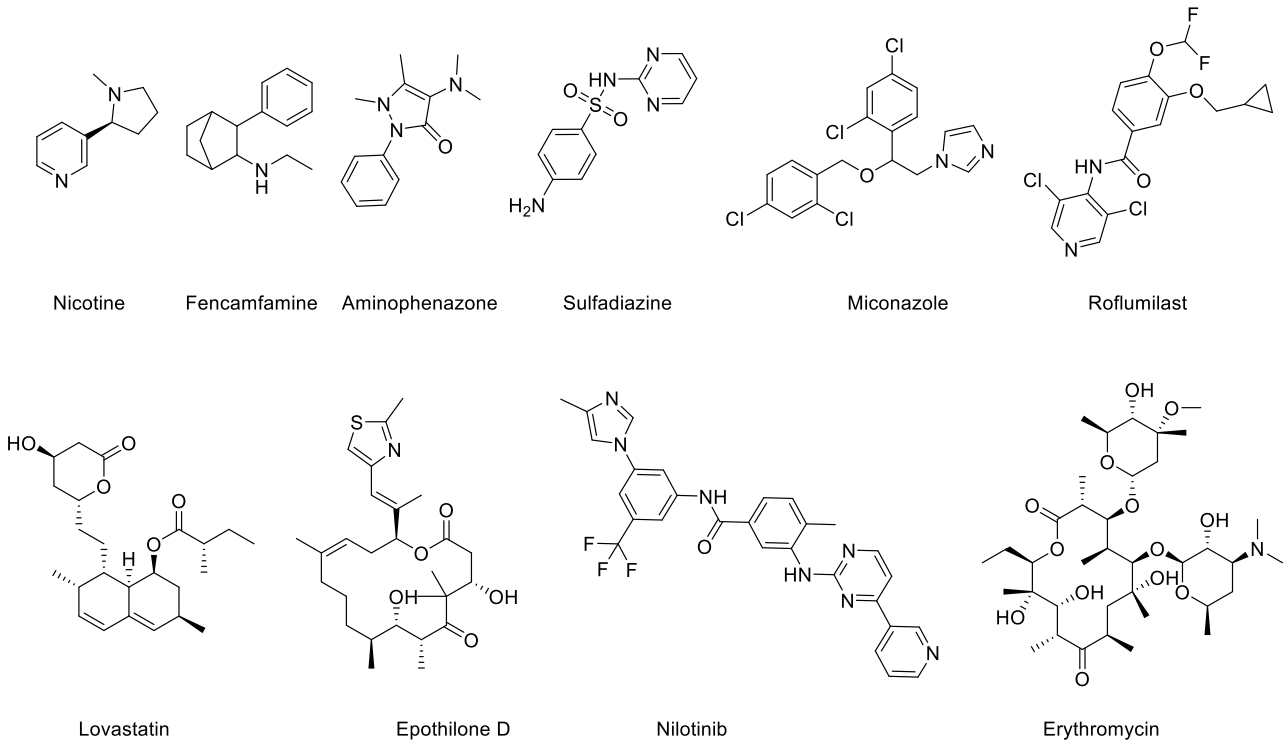

Figure 2. Structure of the 10 drugs used for transfer learning.

Results and Discussion

We trained six LSTMs corresponding to six different primary training sets, namely: 1) 344,319 compounds from ChEMBL containing only molecules reported with high quality datapoints for single protein targets;[18] 2) a random subset of 40,000 molecules from set 1; 3) all compounds up to a size of 50 heavy (non-hydrogen) atoms from DrugBank,[15] which were 5,104 compounds; 4) 40,986 fragments up to 17 atoms collected from various catalogs; 5) 500,000 molecules randomly selected from the fragment database FDB17; 6) sets 1, 3, 4 and 5 combined. As case studies for transfer learning we selected 10 different drug molecules covering a broad range of size and complexity from very small molecules such as nicotine or aminophenazone, to typical drug molecules such as nilotinib and lovastatin, and up to macrocyclic natural products such as epothilone D and erythromycin (Figure 2).

For each of the six LSTMs we performed primary training for 50 epochs using the default learning rate of 0.01 (LSTM6 was trained for 100 epochs considering large training set size, see method for details). We then performed transfer learning for each LSTM using each drug for 20 epochs, using learning rates ranging from 0.0001 to 0.01, generating new molecules after 5, 10, 15 and 20 epochs. We collected all generated SMILES, removed invalid SMILES and duplicates, and applied structural filters to eliminate problematic functional groups. Finally, we selected high-similarity analogs to the drug used for transfer learning using a combined filter considering an Avalon fingerprint[19] Tanimoto cut-off value to constrain substructures, as well as an Xfp[20] city-block-distance cut-off value to constrain overall molecular size, shape, and pharmacophore.

Each LSTM performed differently with each test molecule and learning rate. Analyzing the overall data across all 10 drugs showed that increasing the learning rate of transfer learning led to a strong reduction of the overall number of generated molecules at the level of correct, unique and functionally filtered SMILES. On the other hand, the number of high similarity analogs produced strongly increased with increasing learning rates, such that an optimum was reached at a learning

rate of 0.001 (Figure 3). Training with the large set of ChEMBL molecules produced almost no high similarity analogs (LSTM1, black line), while training all other datasets including the combined set (LSTM2-6) produced comparable numbers of high similarity analogs. It should be noted that the performance of LSTM1 can be tuned by training it for a longer time and tweaking the learning rate.

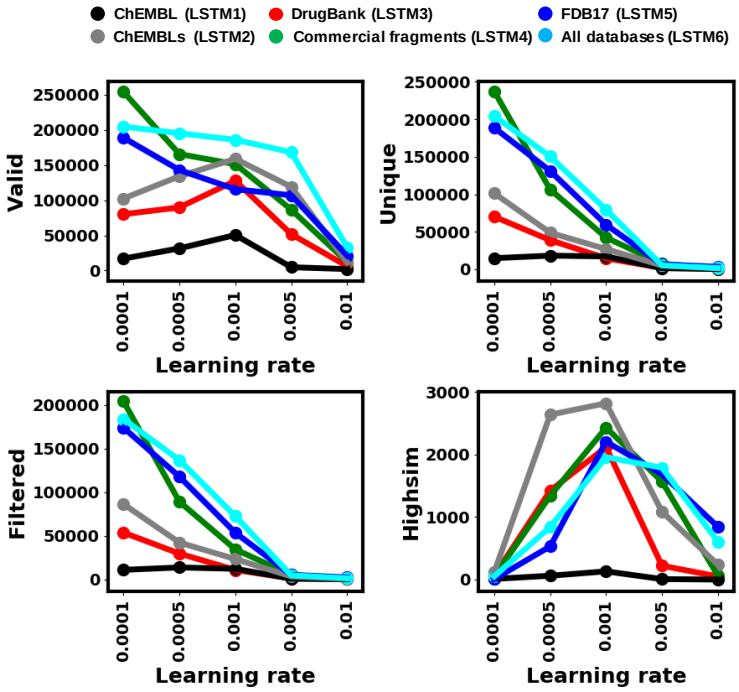

Figure 3. Number of SMILES generated by LSTM1-6 upon transfer learning across all 10 drugs in Figure 2 as a function of learning rates (x-axis). (black line): LSTM1, 344,319 ChEMBL compounds, (grey line): LSTM2, a random subset of 40,000 of the ChEMBL compounds, (red line): LSTM3, 5,104 drugs from DrugBank, (green line): LSTM4, 40,986 commercially available fragments, (blue line): LSTM5, 500,000 fragments from FDB17, (cyan line): LSTM6, datasets 1, 3, 4 and 5 combined. Valid = total number of valid SMILES found across all drugs. For each of the 10 drugs, 200,000 characters were sampled using the respective fine-tuned LSTM model after 5, 10, 15, and 20 epochs. Unique = number of valid SMILES remaining after removing duplicates. Filtered: number of SMILES remaining after removing undesirable functional groups. Highsim: number of SMILES for molecules with Avalon Tanimoto similarity > 0.7 and Xfp city block distance less than Xfp cutoff distance (Xfp cutoff distance = heavy atom count of a drug × 30).

Analyzing the production of analogs for each drug separately showed that LSTM4 and LSTM5, which were trained with relatively small, fragment sized molecules, produced the largest number of analogs among all six LSTMs tested for the large natural product target molecules lovastatin, epothilone D and erythromycin (Figure 4). In these cases, we observed a drift towards the size of the target molecule as the learning rate increased, showing that transfer learning partly consisted in learning how to make these large molecules from fragments.

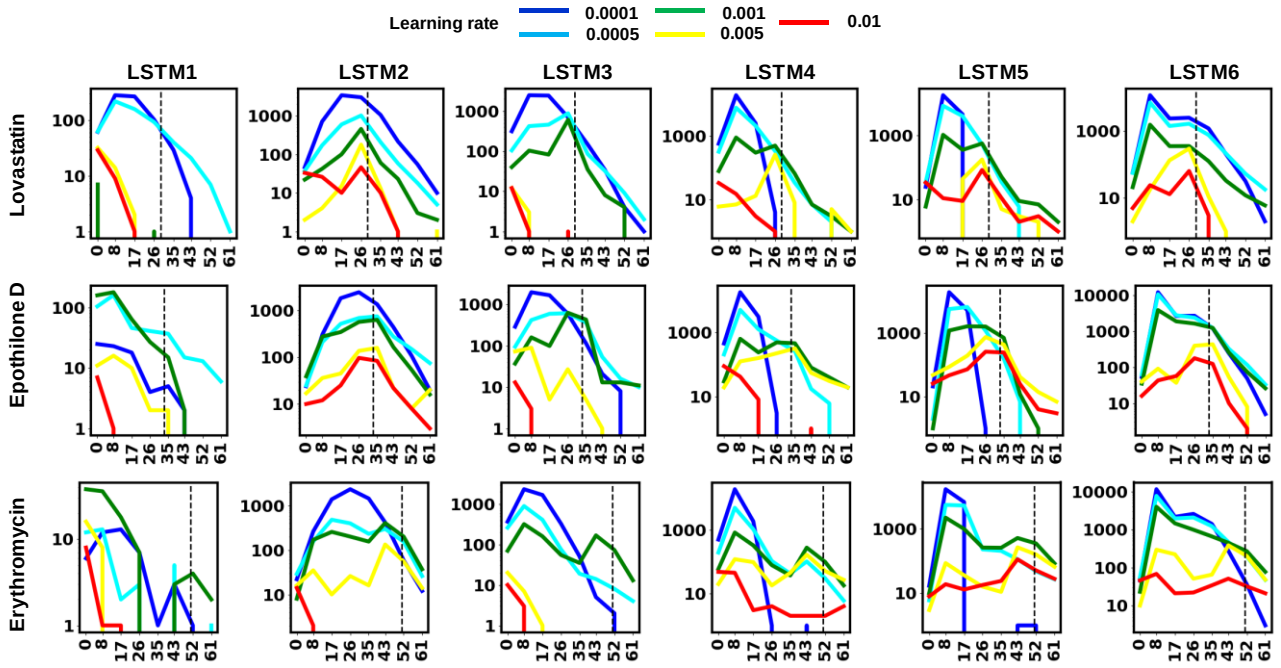

Figure 4. Molecular size (heavy atom count) histogram of generated molecules as a function of learning rate for lovastatin, epothilone D and erythromycin. The vertical dashed line indicates the size of the drug.

To gain a closer insight into the generated molecules, we grouped all compounds generated for each drug by each LSTM across the different learning rates. All LSTMs produced molecules covering the entire Avalon similarity range, spanning from a vast majority of extremely low similarity compounds, to a small fraction of molecules in the high similarity range (Figure 5). A large fraction of these molecules was unique to each LSTM, indicating that the primary training set strongly influenced the molecule generation process (Table 1). Interestingly LSTM6 trained with all sets combined also produced a majority of molecules not generated by any of the other LSTMs. Note that LSTM5 trained with FDB17 did not produce any analogs with halogen containing drugs because FDB17 fragments do not contain any halogens.

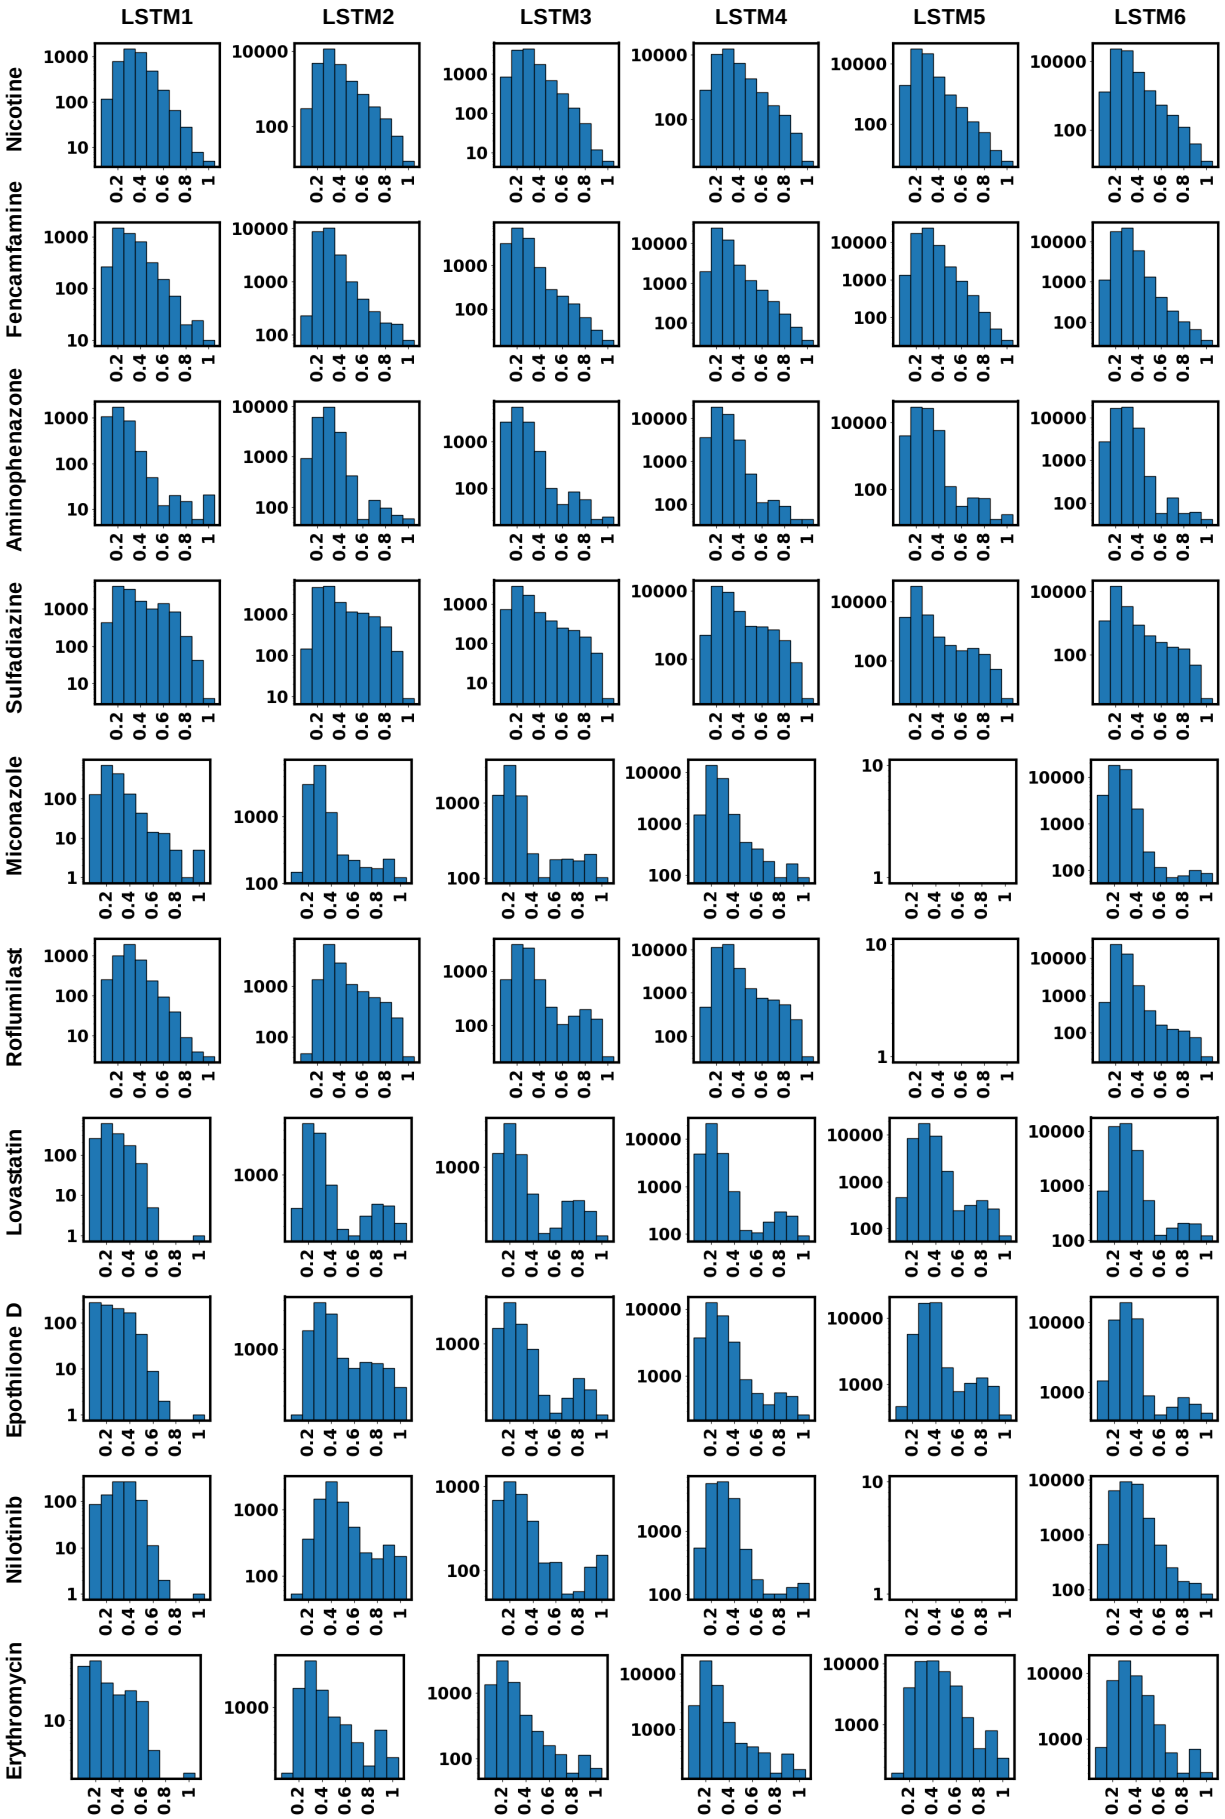

Figure 5. Avalon fingerprint similarity histogram (logarithmmic scale) for all molecules produced by the LSTMs upon transfer learning with the indicated drugs and passing functional group filters.

Table 1. Number of unique/overall filtered compounds produced by the different LSTM neural networks. [a)]

Neural Network	**LSTM1**	**LSTM2**	**LSTM3**	**LSTM4**	**LSTM5**	**LSTM6**	
Source database	ChEMBL	ChEMBLs	DrugBank	Commercial Fragments	FDB17	All databases	Unique across
training cpds.	344,319	40,000	5,104	40,986	500,000	890,409	LSTMs
Nicotine	3008/4389	21K/24K	10K/12K	31K/35K	56K/59K	50K/53K	179K
Fencamfamine	3377/4302	22K/24K	15K/16K	40K/43K	54K/55K	47K/48K	187K
Aminophenazone	2953/3968	19K/20K	10K/11K	35K/38K	64K/65K	42K/44K	179K
Sulfadiazine	11K/12K	14K/15K	5614/6791	26K/28K	42K/43K	20K/21K	123K
Miconazole	1292/1449	10K/11K	6320/6851	25K/25K	0/0	38K/39K	83K
Roflumilast	4036/4268	13K/14K	7550/8183	31K/32K	0/0	39K/40K	97K
Lovastatin	1304/1422	10K/11K	8032/8631	32K/33K	38K/38K	31K/32K	124K
Epothilone D	868/962	11K/11K	7768/8264	29K/30K	45K/45K	46K/46K	142K
Nilotinib	806/862	6941/7216	3384/3636	16K/16K	0/0	27K/28K	55K
Erythromycin	180/219	9235/9535	6764/7116	29K/29K	40K/40K	40K/41K	127K

The pooled set of all filtered molecules for each LSTM had a size distribution close to that of the primary training set, while high similarity analogs covered the range of target molecules (Figure 6a/c). The synthetic accessibility scores[21] of the generated molecules was as expected from the training set, with LSTMs trained with experimental molecules (ChEMBL, DrugBank, commercial fragments) producing molecules with favorable low-value scores, while FDB17 consisting of possible but not synthesized molecules gave less favorable, high-value score, for both filtered and high-similarity analogs (Figure 6c/d).

Most of these high similarity analogs were produced by more than one LSTM, with often less than half of the generated molecules being unique to the LSTM (Table 2). Some of the analogs shared by more than one LSTM were in fact already documented in ChEMBL (Table 3). These known analogs often featured one atom insertions, deletions, substitutions, or inversions (Figure 7). Despite their high similarity to the targets, these analogs were structurally diverse, as shown by the large number of different Bemis-Murcko scaffolds[22] present among these high similarity analogs (Table 4). Note that performing additional molecule generation runs on a single LSTM produced new analogs at each run throughout the Avalon similarity range except for values above 0.9 where only a limited number of compounds are possible. This is illustrated here for the case of nicotine, miconazole and erythromycin analogs produced by LSTM4 trained with commercial fragments,

considering only compounds passing the shape and pharmacophore similarity (Xfp) cut-off (Figure 8 a-c). The range of molecules produced is illustrated here for miconazole by an interactive substructure fingerprint similarity map (Figure 8d).[23] This suggests that many additional such analogs of each drugs can potentially be produced by running LSTMs for longer periods.

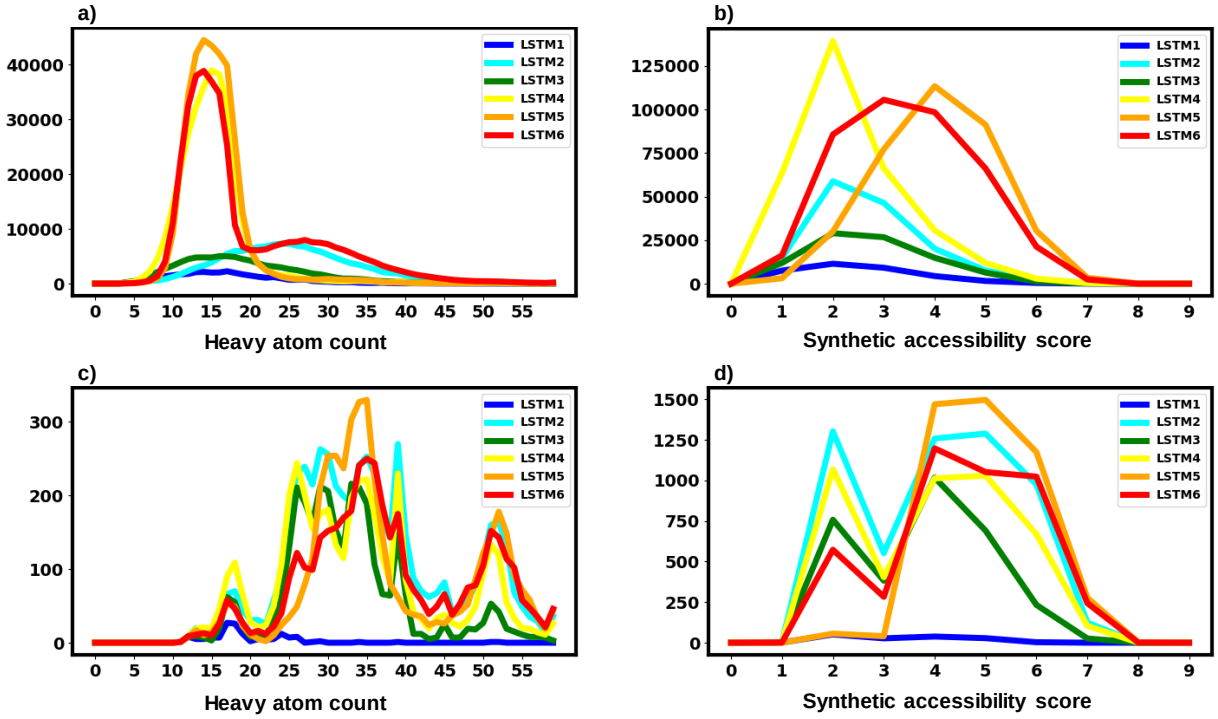

Figure 6. Molecular size histograms and synthetic accessibility score for all filtered (**a-b**) and Highsim (**c-d**) molecules produced by LSTMs summed over the ten different drugs.

Table 2. Number of unique/total high similarity drug analogs produced by the different LSTM neural networks.

Neural Network	**LSTM1**	**LSTM2**	**LSTM3**	**LSTM4**	**LSTM5**	**LSTM6**	
Source database	ChEMBL	ChEMBLs	DrugBank	Commercial Fragments	FDB17	All databases	Unique across LSTMs
training cpds.	344,319	40,000	5,104	40,986	500,000	890,409	
Nicotine	0/23	32/82	1/32	32/93	9/47	16/67	166
Fencamfamine	15/42	126/218	40/96	130/231	92/164	41/114	580
Aminophenazone	5/26	34/96	23/71	38/99	22/66	19/65	223
Sulfadiazine	6/27	19/59	11/37	28/74	8/30	2/25	124
Miconazole	2/10	301/500	268/438	174/336	0/0	153/256	1134
Roflumilast	8/15	319/557	117/283	351/585	0/0	45/166	1126
Lovastatin	0/1	631/986	460/757	352/625	487/728	289/530	2729
Epothilone D	0/1	911/1301	561/831	807/1160	1595/2039	1163/1511	5707
Nilotinib	0/1	506/666	180/321	218/381	0/0	243/355	1362
Erythromycin	0/2	832/1042	174/243	524/709	1243/1444	1105/1288	4190

Table 3. Known bioactive from ChEMBL produced by LSTMs.[a)]

	Known	LSTM1	LSTM2	LSTM3	LSTM4	LSTM5	LSTM6	Unique
Nicotine	41	4	10	5	7	5	6	12
Fencamfamine	12	0	0	0	0	0	0	0
Aminophenazone	134	2	4	3	4	3	4	5
Sulfadiazine	24	0	0	0	2	0	0	2
Miconazole	150	1	17	12	8	0	9	19
Roflumilast	19	1	1	0	0	0	2	2
Lovastatin	49	0	11	12	10	6	11	16
Epothilone D	75	0	5	4	6	4	7	9
Nilotinib	41	0	3	2	1	0	3	4
Erythromycin	201	0	4	1	2	2	1	5

a) Number of known bioactive analogs of each drug found by similarity search in ChEMBL and numbers found in high similarity analogs for each LSTM. Nearest neighbors of each drug were extracted by performing similarity search for a given drug on the ChEMBL website.

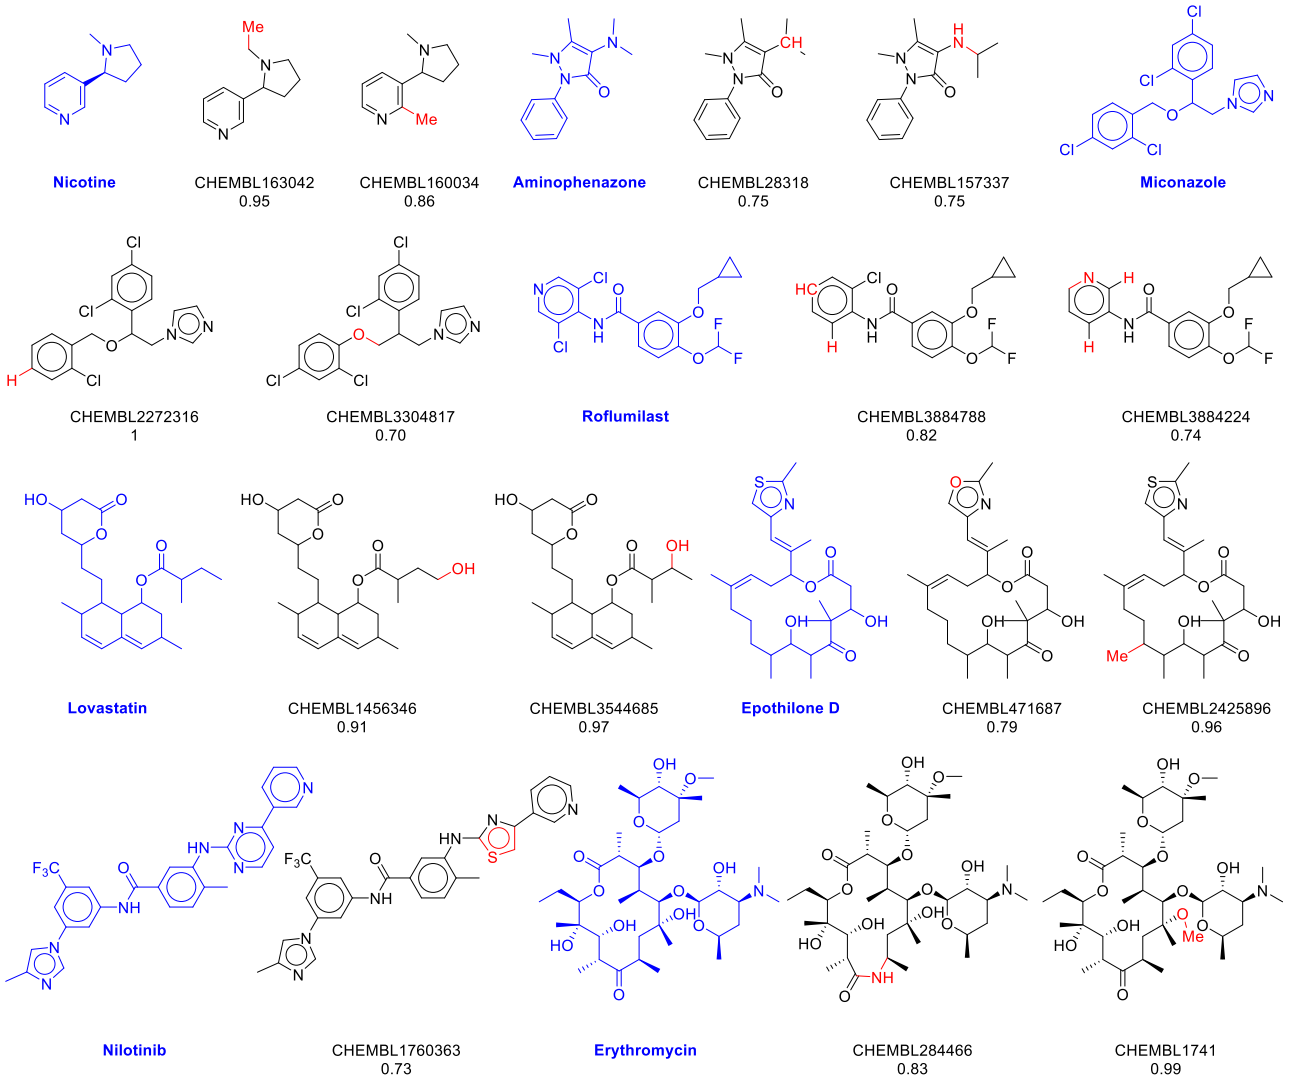

Figure 7. Examples of high similarity drug analogs produced by LSTM and documented in ChEMBL. None of the analogs shown were included in the training sets. The chirality is shown in part of the drugs for clarity of the drawing, however LSTM is run on achiral SMILES and produce analogs without defined stereochemistry.

Table 4. Number of unique/total Bemis-Murcko scaffolds among high similarity drug analogs produced by the different LSTM neural networks.

Neural Network	**LSTM1**	**LSTM2**	**LSTM3**	**LSTM4**	**LSTM5**	**LSTM6**	
Source database	ChEMBL	ChEMBLs	DrugBank	Commercial Fragments	FDB17	All databases	Total Unique
training cpds.	344,319	40,000	5,104	40,986	500,000	890,409	Scaffolds
Nicotine	0/11	4/20	0/16	4/29	0/16	1/23	35
Fencamfamine	7/14	55/92	17/39	41/82	54/90	18/44	247
Aminophenazone	0/7	3/19	3/18	12/30	5/20	3/15	51
Sulfadiazine	1/13	6/21	1/12	5/20	3/15	1/12	32
Miconazole	0/4	65/115	34/78	23/68	0/0	40/76	225
Roflumilast	0/6	75/152	12/75	101/186	0/0	9/49	288
Lovastatin	0/1	143/276	66/175	71/189	113/223	38/108	613
Epothilone D	0/1	491/807	280/538	410/675	781/1145	481/756	2960
Nilotinib	0/1	270/387	109/215	146/278	0/0	91/161	773
Erythromycin	0/1	372/546	87/145	263/413	611/806	510/705	2123

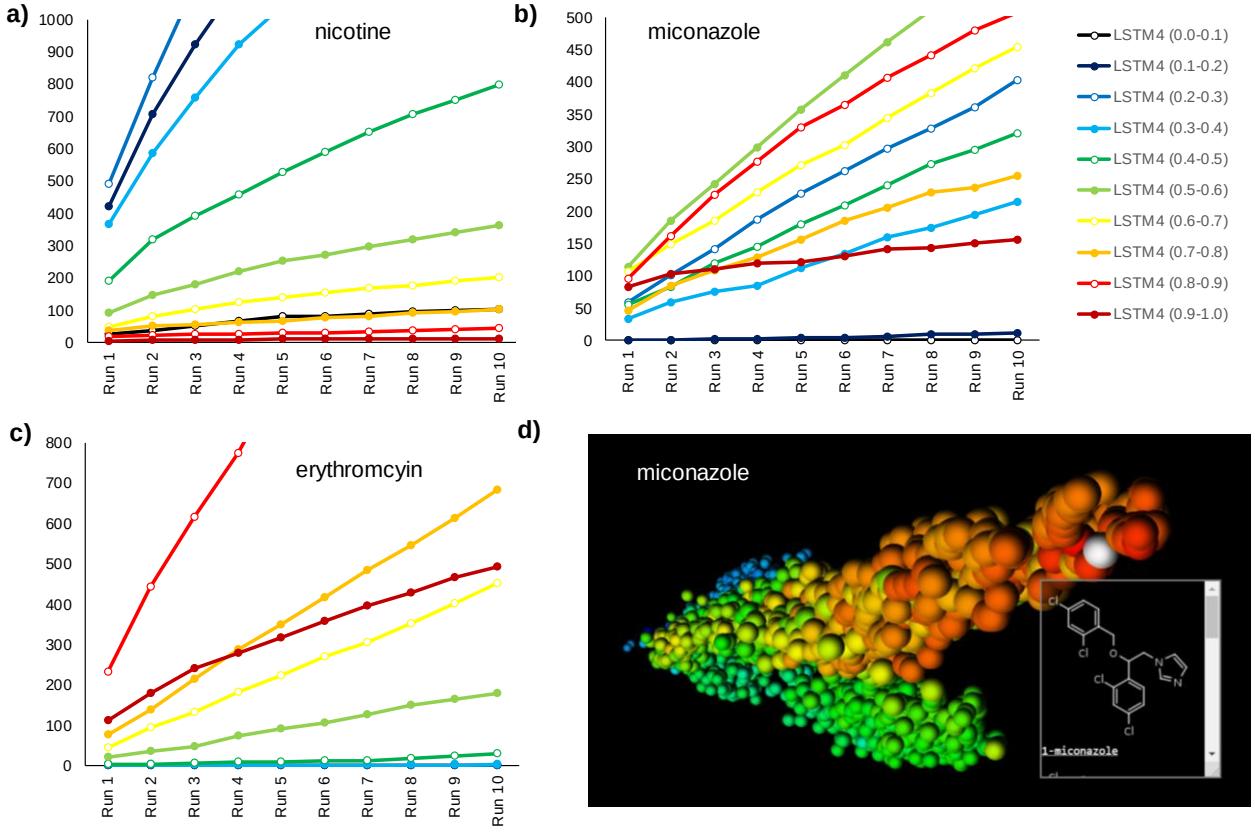

Figure 8. Production of analogs with LSTM4 at 0.001 learning rate upon additional LSTM runs, as function of Avalon Tanimoto similarity to the drug. Only analogs passing the Xfp similarity cut-off were retained. **(a)** cumulative number of unique nicotine analogs upon additional runs as function of Avalon similarity, **(b)** same as (a) for miconazole, **(c)** same as (a) for erythromycin. **(d)** Substructure fingerprint similarity map of miconazole analog produced in (b), color-coded by the Avalon Tanimoto similarity from highest (red) to lowest (blue). The interactive 3D-map is available at the following link: http://gdbtools.unibe.ch:8080/webMolCS/yourSIM.html?jobID=1537890147494&fp=Sfp

Conclusion

Here we trained LSTM generative neural networks with SMILES of drug-like molecules from ChEMBL and DrugBank, or of fragments from commercial catalogs or from FDB17, and performed transfer learning with single drug compounds to generate new analogs of these drugs. We found that LSTMs trained with fragments produced drug analogs as efficiently as LSTMs trained with full size drug-like molecules from ChEMBL or DrugBank. In the case of large natural products such as lovastatin, epothilone D or erythromycin, LSTMs trained with fragments readily learned to assemble large molecules and produced more high similarity analogs of these drugs than LSTMs trained with full-sized molecules, showing that transfer learning informs rules to assemble small fragments into drug-like molecules. Several of the high similarity analogs produced by LSTMs were already documented in ChEMBL and featured one atom changes, however the overall structural diversity of these analogs was high, as attested by a large number of scaffolds. Neural networks trained with approximately 40,000 molecules as simple as a set of commercially available fragments performed excellently in this application, suggesting that fragment-based LSTM neural networks offer a promising method for new molecule generation.

Methods

Compound Databases for LSTM Training. *1) ChEMBL*: The ChEMBL22 database was downloaded from http://www.ebi.ac.uk/chembl. Thereafter, the database was filtered to retain only the compounds reported against a "single protein" target where the source organism was either human or rat, having an activity value (IC_{50}, EC_{50}, EC_{50}, *Ki or* K_D) of ≤10 μM and heavy atom count of ≤50.[18] *2) ChEMBLs:* this set was created by randomly selecting 40,000 compounds from the ChEMBL set mentioned above. *3) DrugBank:* DrugBank database version 5.0.11 was downloaded from http://www.drugbank.ca and filtered to retain only the compounds having ≤50

heavy atoms. *4)* Commercial fragments: fragment like molecules were collected from various suppliers, after which molecules obeying Congreve's rule of three criteria and having heavy atom count of ≤17 were retained in the set. *5) FDB17:* this set was created by randomly sampling 500,000 compounds from entire FDB17 database. *6)* All databases: this set was created by combining the databases *1, 3, 4* and *5.* All molecules were process using the JChem Chemaxon package. Molecules were parsed in non-isomeric unique SMILES format, counter ions were removed from molecules, valence errors were checked, molecules were protonated at pH 7.4, and duplicate molecules were removed in the context of each database. For each database, the plain text file containing the unique SMILES notation of molecules was used as input to train the LSTM model.

LSTM model and Primary Training. All the LSTM models reported herein were constructed using Keras version 2.0.9, a Python based high-level neural network learning library with a TensorFlow-gpu backend. For primary training, we trained six different LSTM models using six different databases namely ChEMBL, ChEMBLs, DrugBank, Commercial fragments, FDB17 and All databases. The architecture of each of these LSTM models contains three LSTM layers (each of size 512), each followed by a Dropout layer of size 512, with dropout-rate of 0.2 to avoid model-overfitting. The output of hidden layer is then processed through the TimeDistributed Layer and Output layer with the softmax as an activation function. All LSTM models were trained using "adagrad" as an optimizer with a learning rate of 0.01 and categorical cross entropy as the loss function. The number of epochs, batch size, sequence length, and vocabulary size for each of these models were: ChEMBL (50, 64, 64, 35), ChEMBLs (50, 32, 64, 35), DrugBank (50, 16, 64, 27), Commercial fragments (50, 32, 54, 34), FDB17 (50, 64, 64, 21) and All databases (100, 64, 64, 40) respectively.

Transfer Learning. Each of the six primary LSTM models were fine-tuned (transfer learning) with respect to each of the 10 drugs mentioned in this paper. For each drug, five independent fine-tuned LSTM models were generated using five different learning rates: 0.0001, 0.0005, 0.001,

0.005 and 0.01. For each model, the plain text file containing the unique SMILES notation of the drug repeated for 20 times was used as the input. Each model was fine-tuned for 20 epochs using a batch size of 4 and a sequence length of 64.

Sampling and Processing of SMILES. For each drug 200,000 characters were sampled using the respective fine-tuned LSTM model after 5, 10, 15, and 20 epochs. The Numpy *random.choice* method was used to select the character given the predicted probabilities over the vocabulary. After sampling, the generated SMILES (newline character was used as delimiter to separate the SMILES of different molecules) were processed using the RDkit library. Molecules which were successfully processed by RDkit were considered as valid molecules. Thereafter, molecules were protonated at pH 7.4 using the JChem Chemaxon library; duplicate molecules and molecules containing unstable functional groups were also removed from the list.

Avalon and Xfp Fingerprints. An Avalon substructure fingerprint containing 1,024 bits was computed using RDkit and the Avalon toolkit. The Xfp topological pharmacophore and shape fingerprint was computed using in-housed written Java-program using the Jchem Chemaxon library as a starting point. Similarities between molecules were quantified using the Tanimoto coefficient and the City block distance, respectively for Avalon and Xfp fingerprints.

Authors' contributions. MA designed and realized the study and wrote the paper. FS, NS and JLR co-designed and supervised the study and wrote the paper.

Notes. The authors declare no competing financial interest.

Acknowledgments. This work was supported financially by a grant of NIBR to MA.

References

1. Bleicher, K. H.; Bohm, H. J.; Muller, K.; Alanine, A. I. Hit and Lead Generation: Beyond High-Throughput Screening. *Nat. Rev. Drug Discovery* **2003,** *2*, 369-378.

2. Renner, S.; Popov, M.; Schuffenhauer, A.; Roth, H. J.; Breitenstein, W.; Marzinzik, A.; Lewis, I.; Krastel, P.; Nigsch, F.; Jenkins, J.; Jacoby, E. Recent Trends and Observations in the Design of High-Quality Screening Collections. *Future Med. Chem.* **2011,** *3*, 751-766.

3. Reymond, J. L.; Ruddigkeit, L.; Blum, L. C.; Van Deursen, R. The Enumeration of Chemical Space. *WIREs comput. Mol. Sci.* **2012**, doi: 10.1002/wcms.1104.

4. Awale, M.; Visini, R.; Probst, D.; Arus-Pous, J.; Reymond, J. L. Chemical Space: Big Data Challenge for Molecular Diversity. *Chimia* **2017,** *71*, 661-666.

5. Scior, T.; Bender, A.; Tresadern, G.; Medina-Franco, J. L.; Martinez-Mayorga, K.; Langer, T.; Cuanalo-Contreras, K.; Agrafiotis, D. K. Recognizing Pitfalls in Virtual Screening: A Critical Review. *J. Chem. Inf. Model.* **2012,** *52*, 867-881.

6. Schmidhuber, J. Deep Learning in Neural Networks: An Overview. *Neural Networks* **2015,** *61*, 85-117.

7. Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G. S.; Davis, A.; Dean, J.; Devin, M.; Ghemawat, S.; Goodfellow, I.; Harp, A.; Irving, G.; Isard, M.; Jia, Y.; Jozefowicz, R.; Kaiser, L.; Kudlur, M.; Levenberg, J.; Mané, D.; Monga, R.; Moore, S.; Murray, D.; Olah, C.; Schuster, M.; Shlens, J.; Steiner, B.; Sutskever, I.; Talwar, K.; Tucker, P.; Vanhoucke, V.; Vasudevan, V.; Viégas, F.; Vinyals, O.; Warden, P.; Wattenberg, M.; Wicke, M.; Yu, Y.; Zheng, X.; Research, G. Tensorflow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. *arXiv:1603.04467* **2016**.

8. 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.

9. 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*, D1083-D1090.

10. Segler, M. H. S.; Kogej, T.; Tyrchan, C.; Waller, M. P. Generating Focussed Molecule Libraries for Drug Discovery with Recurrent Neural Networks. *arXiv:1701.01329* **2017**.

11. Merk, D.; Friedrich, L.; Grisoni, F.; Schneider, G. De Novo Design of Bioactive Small Molecules by Artificial Intelligence. *Mol. Inform.* **2018,** *37*, 1700153.

12. Ertl, P.; Lewis, R.; Martin, E.; Polyakov, V. In Silico Generation of Novel, Drug-Like Chemical Matter Using the Lstm Neural Network. *arXiv:1712.07449v2* **2017**.

13. Olivecrona, M.; Blaschke, T.; Engkvist, O.; Chen, H. Molecular De-Novo Design through Deep Reinforcement Learning. *J. Cheminform.* **2017,** *9*, 48.

14. Li, Y.; Zhang, L.; Liu, Z. Multi-Objective De Novo Drug Design with Conditional Graph Generative Model. *J. Cheminform.* **2018,** *10*, 33.

15. Law, V.; Knox, C.; Djoumbou, Y.; Jewison, T.; Guo, A. C.; Liu, Y.; Maciejewski, A.; Arndt, D.; Wilson, M.; Neveu, V.; Tang, A.; Gabriel, G.; Ly, C.; Adamjee, S.; Dame, Z. T.; Han, B.; Zhou, Y.; Wishart, D. S. Drugbank 4.0: Shedding New Light on Drug Metabolism. *Nucleic Acids Res.* **2014,** *42*, D1091-D1097.

16. Visini, R.; Awale, M.; Reymond, J.-L. Fragment Database Fdb-17. *J. Chem. Inf. Model.* **2017,** *57*, 700-709.

17. Ruddigkeit, L.; van Deursen, R.; Blum, L. C.; Reymond, J. L. Enumeration of 166 Billion Organic Small Molecules in the Chemical Universe Database Gdb-17. *J. Chem. Inf. Model.* **2012,** *52*, 2864-2875.

18. Awale, M.; Reymond, J. L. The Polypharmacology Browser Ppb2: Target Prediction Combining Nearest Neighbors with Machine Learning. *ChemRXiv* **2018**, https://doi.org/10.26434/chemrxiv.6895646.v1.

19. Gedeck, P.; Rohde, B.; Bartels, C. Qsar--How Good Is It in Practice? Comparison of Descriptor Sets on an Unbiased Cross Section of Corporate Data Sets. *J. Chem. Inf. Model.* **2006,** *46*, 1924-1936.

20. Awale, M.; Reymond, J. L. Atom Pair 2d-Fingerprints Perceive 3d-Molecular Shape and Pharmacophores for Very Fast Virtual Screening of Zinc and Gdb-17. *J. Chem. Inf. Model.* **2014,** *54*, 1892-1897.

21. Ertl, P.; Schuffenhauer, A. Estimation of Synthetic Accessibility Score of Drug-Like Molecules Based on Molecular Complexity and Fragment Contributions. *J. Cheminform.* **2009,** *1*, 8.

22. Bemis, G. W.; Murcko, M. A. The Properties of Known Drugs. 1. Molecular Frameworks. *J. Med. Chem.* **1996,** *39*, 2887-2893.

23. Awale, M.; Probst, D.; Reymond, J. L. Webmolcs: A Web-Based Interface for Visualizing Molecules in Three-Dimensional Chemical Spaces. *J. Chem. Inf. Model.* **2017,** *57*, 643-649.

Graphics for the Table of Contents: