

LEARNING TO SMILE(S)

Stanisław Jastrzebski, Damian Leśniak & Wojciech Marian Czarnecki

Faculty of Mathematics and Computer Science

Jagiellonian University

Kraków, Poland

stanislaw.jastrzebski@uj.edu.pl

ABSTRACT

This paper shows how one can directly apply natural language processing (NLP) methods to classification problems in cheminformatics. Connection between these seemingly separate fields is shown by considering standard textual representation of compound, SMILES. The problem of activity prediction against a target protein is considered, which is a crucial part of computer aided drug design process. Conducted experiments show that this way one can not only outrank state of the art results of hand crafted representations but also gets direct structural insights into the way decisions are made.

1 INTRODUCTION

Computer aided drug design has become a very popular technique for speeding up the process of finding new biologically active compounds by drastically reducing number of compounds to be tested in laboratory. Crucial part of this process is virtual screening, where one considers a set of molecules and predicts whether the molecules will bind to a given protein. This research focuses on ligand-based virtual screening, where the problem is modelled as a supervised, binary classification task using only knowledge about ligands (drug candidates) rather than using information about the target (protein).

One of the most underrepresented application areas of deep learning (DL) is believed to be cheminformatics (Unterthiner et al., 2014; Bengio et al., 2012), mostly due the fact that data is naturally represented as graphs and there are little direct ways of applying DL in such setting (Henaff et al., 2015). Notable examples of DL successes in this domain are winning entry to Merck competition in 2012 (Dahl et al., 2014) and Convolutional Neural Network (CNN) used for improving data representation (Duvenaud et al., 2015). To the authors best knowledge all of the above methods use hand crafted representations (called fingerprints) or use DL methods in a limited fashion. The main contribution of the paper is showing that one can directly apply DL methods (*without* any customization) to the textual representation of compound (where characters are atoms and bonds). This is analogous to recent work showing that state of the art performance in language modelling can be achieved considering character-level representation of text (Kim et al., 2015; Jozefowicz et al., 2016).

1.1 REPRESENTING MOLECULES

Standard way of representing compound in any chemical database is called SMILES, which is just a string of atoms and bonds constructing the molecule (see Fig. 1) using a specific walk over the graph. Quite surprisingly, this representation is rarely used as a base of machine learning (ML) methods (Worachartcheewan et al., 2014; Toropov et al., 2010).

Most of the classical ML models used in cheminformatics (such as Support Vector Machines or Random Forest) work with constant size vector representation through some predefined embedding (called *fingerprints*). As a result many such fingerprints have been proposed across the years (Hall & Kier, 1995; Steinbeck et al., 2003). One of the most common ones are the substructural ones - analogous of bag of word representation in NLP, where fingerprint is defined as a set of graph *templates* (SMARTS), which are then matched against the molecule to produce binary (set of words)

or count (bag of words) representation. One could ask if this is really necessary, having at one's disposal DL methods of feature learning.

1.2 ANALOGY TO SENTIMENT ANALYSIS

The main contribution of this paper is identifying analogy to NLP and specifically sentiment analysis, which is tested by applying state of the art methods (Mesnil et al., 2014) directly to SMILES representation. The analogy is motivated by two facts. First, small local changes to structure can imply large overall activity change (see Fig. 2), just like sentiment is a function of sentiments of different clauses and their connections, which is the main argument for effectiveness of DL methods in this task (Socher et al., 2013). Second, perhaps surprisingly, compound graph is almost always nearly a tree. To confirm this claim we calculate molecules diameters, defined as a maximum over all atoms of minimum distance between given atom and the longest carbon chain in the molecule. It appears that in practise analyzed molecules have diameter between 1 and 6 with mean 4. Similarly, despite the way people write down text, human thoughts are not linear, and sentences can have complex clauses. Concluding, in organic chemistry one can make an analogy between longest carbon chain and sentence, where branches stemming out of the longest chain are treated as clauses in NLP.




Figure 1: SMILES produced for the compound in the figure is N(c1)ccc1N.




Figure 2: Substituting highlighted carbon atom with nitrogen renders compound inactive.




Figure 3: Visualization of CNN filters of size 5 for active (top row) and inactive molecules.

2 EXPERIMENTS

Five datasets are considered. Except SMILES, two baseline fingerprint compound representations are used, namely MACCS Ewing et al. (2006) and Klekota–Roth Klekota & Roth (2008) (KR; considered state of the art in substructural representation (Czarnecki et al., 2015)). Each dataset is fairly small (mean size is 3000) and most of the datasets are slightly imbalanced (with mean class ratio around 1:2). It is worth noting that chemical databases are usually fairly big (ChEMBL size is 1.5M compounds), which hints at possible gains by using semi-supervised learning techniques.

Tested models include both traditional classifiers: Support Vector Machine (SVM) using Jaccard kernel, Naive Bayes (NB), Random Forest (RF) as well as neural network models: Recurrent Neural Network Language Model (Mikolov et al., 2011b) (RNNLM), Recurrent Neural Network (RNN) many to one classifier, Convolutional Neural Network (CNN) and Feed Forward Neural Network with ReLU activation. Models were selected to fit two criteria: span state of the art models in single target virtual screening (Czarnecki et al., 2015; Smusz et al., 2013) and also cover state of the art models in sentiment analysis. For CNN and RNN a form of data augmentation is used, where for each molecule random SMILES walks are computed and predictions are averaged (not doing so degrades strongly performance, mostly due to overfitting). For methods which are not designed to work on string representation (such as SVM, NB, RF, etc.) SMILES are embedded as n-gram models with simple tokenization ([Na+] becomes a single token). For all the remaining ones, SMILES are treated as strings composed of 2-chars symbols (thus capturing atom and its relation to the next one).

Using RNNLM, $p(\text{compound}|\text{active})$ and $p(\text{compound}|\text{inactive})$ are modelled separately and classification is done through logistic regression fitted on top. For CNN, purely supervised version of CONTEXT, current state of the art in sentiment analysis (Johnson & Zhang, 2015), is used. Notable feature of the model is working directly on one-hot representation of the data. Each model is evaluated using 5-fold stratified cross validation. Internal 5-fold grid is used for fitting hyperparameters (truncated in the case of deep models). We use log loss as an evaluation metric to include both

Table 1: Log-loss (\pm std) of each model for a given protein and representation.

	model	5-HT _{1A}	5-HT _{2A}	5-HT ₇	H1	SERT
SMILES	CNN	0.249 \pm 0.015	0.284 \pm 0.026	0.289 \pm 0.041	0.182 \pm 0.030	0.221 \pm 0.032
	SVM	0.255 \pm 0.009	0.309 \pm 0.027	0.302 \pm 0.033	0.202 \pm 0.037	0.226 \pm 0.015
	GRU	0.274 \pm 0.016	0.340 \pm 0.035	0.347 \pm 0.045	0.222 \pm 0.042	0.269 \pm 0.032
	RNNLM	0.363 \pm 0.020	0.431 \pm 0.025	0.486 \pm 0.065	0.283 \pm 0.066	0.346 \pm 0.102
KRFP	SVM	0.262 \pm 0.016	0.311 \pm 0.021	0.326 \pm 0.035	0.188 \pm 0.022	0.226 \pm 0.014
	RF	0.264 \pm 0.029	0.297 \pm 0.012	0.322 \pm 0.038	0.210 \pm 0.015	0.228 \pm 0.022
	NN	0.285 \pm 0.026	0.331 \pm 0.015	0.375 \pm 0.072	0.232 \pm 0.034	0.240 \pm 0.024
	NB	0.634 \pm 0.045	0.788 \pm 0.073	1.201 \pm 0.315	0.986 \pm 0.331	0.726 \pm 0.066
MACCS	SVM	0.310 \pm 0.012	0.339 \pm 0.017	0.382 \pm 0.019	0.237 \pm 0.027	0.280 \pm 0.030
	RF	0.261 \pm 0.008	0.294 \pm 0.015	0.335 \pm 0.034	0.202 \pm 0.004	0.237 \pm 0.029
	NN	0.377 \pm 0.005	0.422 \pm 0.025	0.463 \pm 0.047	0.278 \pm 0.027	0.369 \pm 0.020
	NB	0.542 \pm 0.043	0.565 \pm 0.014	0.660 \pm 0.050	0.477 \pm 0.042	0.575 \pm 0.017

classification results as well as uncertainty measure provided by models. Similar conclusions are true for accuracy.

2.1 RESULTS

Results are presented in Table 1. First, simple n-gram models (SVM, RF) performance is close to hand crafted state of the art representation, which suggests that potentially *any* NLP classifier working on n-gram representation might be applicable. Maybe even more interestingly, current state of the art model for sentiment analysis - CNN - despite small dataset size, outperforms (however by a small margin) traditional models.

Hyperparameters selected for CNN (CONTEXT) are similar to the parameters reported in (Johnson & Zhang, 2015). Especially the maximum pooling (as opposed to average pooling) and moderately sized regions (5 and 3) performed best (see Fig. 3). This effect for NLP is strongly correlated with the fact that small portion of sentence can contribute strongly to overall sentiment, thus confirming claimed molecule-sentiment analogy.

RNN classifier’s low performance can be attributed to small dataset sizes, as commonly RNN are applied to significantly larger volumes of data (Mikolov et al., 2011a). One alternative is to consider semi-supervised version of RNN (Dai & Le, 2015). Another problem is that compound activity prediction requires remembering very long interactions, especially that neighbouring atoms in SMILES walk are often disconnected in the original molecule.

3 CONCLUSIONS

This work focuses on the problem of compounds activity prediction without hand crafted features used to represent complex molecules. Presented analogies with NLP problems, and in particular sentiment analysis, followed by experiments performed with the use of state of the art methods from both NLP and cheminformatics seem to confirm that one can actually learn directly from raw string representation of SMILES instead of currently used embedding. In particular, performed experiments show that despite being trained on relatively small datasets, CNN based solution can actually outperform state of the art methods based on structural fingerprints in ligand-based virtual screening task. At the same time it gives possibility to easily incorporate unsupervised and semi-supervised techniques into the models, making use of huge databases of chemical compounds. It appears that cheminformatics can strongly benefit from NLP and further research in this direction should be conducted.

ACKNOWLEDGMENTS

First author was supported by Grant No. DI 2014/016644 from Ministry of Science and Higher Education, Poland.

REFERENCES

- Yoshua Bengio, Aaron Courville, and Pascal Vincent. Unsupervised feature learning and deep learning: A review and new perspectives. *CoRR*, abs/1206.5538, 2012. URL <http://arxiv.org/abs/1206.5538>.
- Wojciech Marian Czarnecki, Sabina Podlowska, and Andrzej Bojarski. Robust optimization of svm hyperparameters in the classification of bioactive compounds. *Journal of Cheminformatics*, 7(38), 2015.
- George Dahl, Navdeep Jaitly, and Ruslan Salakhutdinov. Multi-task neural networks for QSAR predictions. *CoRR*, abs/1406.1231, 2014. URL <http://arxiv.org/abs/1406.1231>.
- Andrew Dai and Quoc Viet Le. Semi-supervised sequence learning. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 28*, pp. 3061–3069. Curran Associates, Inc., 2015. URL <http://papers.nips.cc/paper/5949-semi-supervised-sequence-learning.pdf>.
- David Kristjanson Duvenaud, Dougal Maclaurin, Jorge Aguilera-Iparraguirre, Rafael Gómez-Bombarelli, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan Prescott Adams. Convolutional networks on graphs for learning molecular fingerprints. *CoRR*, abs/1509.09292, 2015. URL <http://arxiv.org/abs/1509.09292>.
- Todd Ewing, J. Christian Baber, and Miklos Feher. Novel 2d fingerprints for ligand-based virtual screening. *Journal of Chemical Information and Modeling*, 46(6):2423–2431, 2006. URL <http://dx.doi.org/10.1021/ci060155b>.
- Lowell Hall and Lemont Kier. Electrotological state indices for atom types: A novel combination of electronic, topological, and valence state information. *Journal of Chemical Information and Computer Sciences*, 35(6):1039–1045, 1995. URL <http://dblp.uni-trier.de/db/journals/jcisd/jcisd35.html#HallK95>.
- Mikael Henaff, Joan Bruna, and Yann LeCun. Deep convolutional networks on graph-structured data. *CoRR*, abs/1506.05163, 2015. URL <http://arxiv.org/abs/1506.05163>.
- Rie Johnson and Tong Zhang. Effective use of word order for text categorization with convolutional neural networks. In *NAACL HLT 2015, The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 103–112, 2015. URL <http://aclweb.org/anthology/N/N15/N15-1011.pdf>.
- Rafal Jozefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu. Exploring the limits of language modeling. volume abs/1602.02410, 2016. URL <http://arxiv.org/abs/1602.02410>.
- Yoon Kim, Yacine Jernite, David Sontag, and Alexander Rush. Character-aware neural language models. *CoRR*, abs/1508.06615, 2015. URL <http://arxiv.org/abs/1508.06615>.
- Justin Klekota and Frederick Roth. Chemical substructures that enrich for biological activity. *Bioinformatics*, 24(21):2518–2525, 2008. URL <http://dblp.uni-trier.de/db/journals/bioinformatics/bioinformatics24.html#KlekotaR08>.
- Grégoire Mesnil, Tomas Mikolov, Marc’ Aurelio Ranzato, and Yoshua Bengio. Ensemble of generative and discriminative techniques for sentiment analysis of movie reviews. *CoRR*, abs/1412.5335, 2014. URL <http://arxiv.org/abs/1412.5335>.
- Tomas Mikolov, Anoop Deoras, Daniel Povey, Lukás Burget, and Jan Cernocký. Strategies for training large scale neural network language models. In David Nahamoo and Michael Picheny (eds.), *2011 IEEE Workshop on Automatic Speech Recognition & Understanding, ASRU*, pp. 196–201. IEEE, 2011a. URL <http://dx.doi.org/10.1109/ASRU.2011.6163930>.
- Tomas Mikolov, Stefan Kombrink, Anoop Deoras, Lukar Burget, and Jan Cernocky. Rnnlm-recurrent neural network language modeling toolkit. *Proc. of the 2011 ASRU Workshop*, pp. 196–201, 2011b.

- Sabina Smusz, Rafał Kurczab, and Andrzej Bojarski. A multidimensional analysis of machine learning methods performance in the classification of bioactive compounds. *Chemometrics and Intelligent Laboratory Systems*, 128:89–100, 2013.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1631–1642, 2013. URL <http://www.aclweb.org/anthology-new/D/D13/D13-1170.bib>.
- Christoph Steinbeck, Yongquan Han, Stefan Kuhn, Oliver Horlacher, Edgar Luttmann, and Egon Willighagen. The chemistry development kit (cdk): An open-source java library for chemo- and bioinformatics. *Journal of Chemical Information and Computer Sciences*, 43(2):493–500, 2003.
- Andrey Toropov, Alla Toropova, and E Benfenati. Smiles-based optimal descriptors: Qsar modeling of carcinogenicity by balance of correlations with ideal slopes. *European journal of medicinal chemistry*, 45(9):3581–3587, September 2010. URL <http://dx.doi.org/10.1016/j.ejmech.2010.05.002>.
- Thomas Unterthiner, Andreas Mayr, Günter Klambauer, Marvin Steijaert, Jörg Wenger, Hugo Ceulemans, and Sepp Hochreiter. Deep learning as an opportunity in virtual screening. *Deep Learning and Representation Learning Workshop (NIPS 2014)*, 2014.
- Apilak Worachartcheewan, Prasit Mandi, Virapong Prachayasittikul, Alla Toropova, Andrey Toropov, and Chanin Nantasenamat. Large-scale qsar study of aromatase inhibitors using smiles-based descriptors. *Chemometrics and Intelligent Laboratory Systems*, 138(Complete):120–126, 2014.