

Transformer-CNN: Fast and Reliable Tool for QSAR

Pavel Karpov, Guillaume Godin, Igor Tetko

Submitted date: 09/10/2019 • Posted date: 21/10/2019

Licence: CC BY-NC-ND 4.0

Citation information: Karpov, Pavel; Godin, Guillaume; Tetko, Igor (2019): Transformer-CNN: Fast and Reliable Tool for QSAR. ChemRxiv. Preprint.

We present SMILES-embeddings derived from internal encoder state of a Transformer model trained to canonize SMILES as a Seq2Seq problem. Using CharNN architecture upon the embeddings results in a higher quality QSAR/QSPR models on diverse benchmark datasets including regression and classification tasks. The proposed Transformer-CNN method uses SMILES augmentation for training and inference, and thus the prognosis grounds on an internal consensus. Both the augmentation and transfer learning based on embedding allows the method to provide good results for small datasets. We discuss the reasons for such effectiveness and draft future directions for the development of the method. The source code and the embeddings are available on <https://github.com/bigchem/transformer-cnn>, whereas the OCHEM environment (<https://ochem.eu>) hosts its on-line implementation.

File list (1)

Trans-CNN.docx (478.64 KiB)

[view on ChemRxiv](#) • [download file](#)

Transformer-CNN: Fast and Reliable tool for QSAR

Pavel Karpov,^{1,2} Guillaume Godin³ and Igor V. Tetko^{1,2}

¹Helmholtz Zentrum München - Research Center for Environmental Health (GmbH), Institute of Structural Biology, Ingolstädter Landstraße 1, D-85764 Neuherberg, Germany,

²BigChem GmbH, Ingolstädter Landstraße 1, b. 60w, D-85764 Neuherberg, Germany
pavel.karpov and i.tetko@helmholtz-muenchen.de

³Firmenich International SA, Digital Lab, Switzerland, Geneva, Lausanne
guillaume.godin@firmenich.com

Corresponding author: pavel.karpov@helmholtz-muenchen.de

Abstract

We present SMILES-embeddings derived from internal encoder state of a Transformer[1] model trained to canonize SMILES as a Seq2Seq problem. Using CharNN[2] architecture upon the embeddings results in a higher quality QSAR/QSPR models on diverse benchmark datasets including regression and classification tasks. The proposed Transformer-CNN method uses SMILES augmentation for training and inference, and thus the prognosis grounds on an internal consensus. Both the augmentation and transfer learning based on embedding allows the method to provide good results for small datasets. We discuss the reasons for such effectiveness and draft future directions for the development of the method. The source code and the embeddings are available on <https://github.com/bigchem/transformer-cnn>, whereas the OCHEM[3] environment (<https://ochem.eu>) hosts its on-line implementation.

Keywords: Transformer model, Convolutional neural neural networks, Augmentation, QSAR, SMILES, Embeddings, Character-based models, Cheminformatics, Regression, Classification.

Introduction

Quantitative Structure-Activity (Property) Relationship (QSAR/QSPR) approaches find a nonlinear function, often modeled as an artificial neural network (ANN), that estimates the activity/property based on a chemical structure. In the past, the most QSAR works heavily relied on descriptors[4] that represent in a numerical way some features of a complex graph structure of a compound. Amongst numerous families of descriptors, the fragment descriptors that count occurrences of a subgraph in a molecule graph, hold a distinctive status due to simplicity in the calculation and moreover, there is a theoretical proof that one can successfully build any QSAR model with them[5]. Even a small database of compounds contains thousands of fragmental

descriptors and some feature selection algorithm was used to find a proper subset of descriptors for better quality and speed up of the whole modeling process. Thus, feature selection in conjunction with a suitable machine learning method was a key to success[6]. Rise of deep learning[7] allows getting rid of tiresome expert and domain-wise feature construction by delegating this task to a neural network that can extract the most valuable traits of the raw input data required for modeling problem at hand[8, 9].

In this setting, the whole molecule as a SMILES-strings [10, 11] (Simplified Molecular Input Line Entry System) or a graph[12, 13] serves as the input to the neural network. SMILES notation allows writing any complex formula of an organic compound in a string facilitating storage and retrieval information about molecules in databases[14]. It contains all information about the compound sufficient to derive the entire configuration (3D-structure) and has a direct connection to the nature of fragmental descriptors, Fig. 1, thus, making SMILES one of the best representation for QSAR studies.




Fig. 1. Benzylpenicillin canonical SMILES at the top, 2D and 3D structures derived from SMILES with OpenBabel[15] in the middle, and three non-canonical SMILES examples at the bottom. A substructure of the phenyl ring is written in bold font.

One of the first work exploiting direct SMILES input as descriptors used fragmentation of strings into groups of overlapping substrings forming a SMILES-like set or a hologram of a molecule[16]. Within this approach, there was no need to derive a 2D/3D configuration of the molecule with subsequent calculation of descriptors keeping the quality of the models at the same level as with classical descriptors or even better.

In the first place, SMILES strings are sequences of characters; therefore, they can be

analyzed by machine-learning methods suitable for text processing, namely with convolutional and recurrent neural networks. After the demonstration of text understanding from character-level inputs[17], the technique was adopted in chemoinformatics[11, 18–21]. Recently we showed that augmentation of SMILES (using canonical as well as non-canonical SMILES during model training and inference) increases the performance of convolutional models for regression and classification tasks[22].

Technically modern machine-learning models consist of two parts working together. The first one encodes the input data and extracts the most robust features by applying convolutional filters with different receptive fields (RF) or recurrent layers, whereas the second part directly builds the regular model based on these features using standard dense layers as building blocks (so called classical “MLP”), Fig. 2. Though powerful convolutional layers can effectively encode the input to its internal representation, usually one needs a considerable training dataset and computational

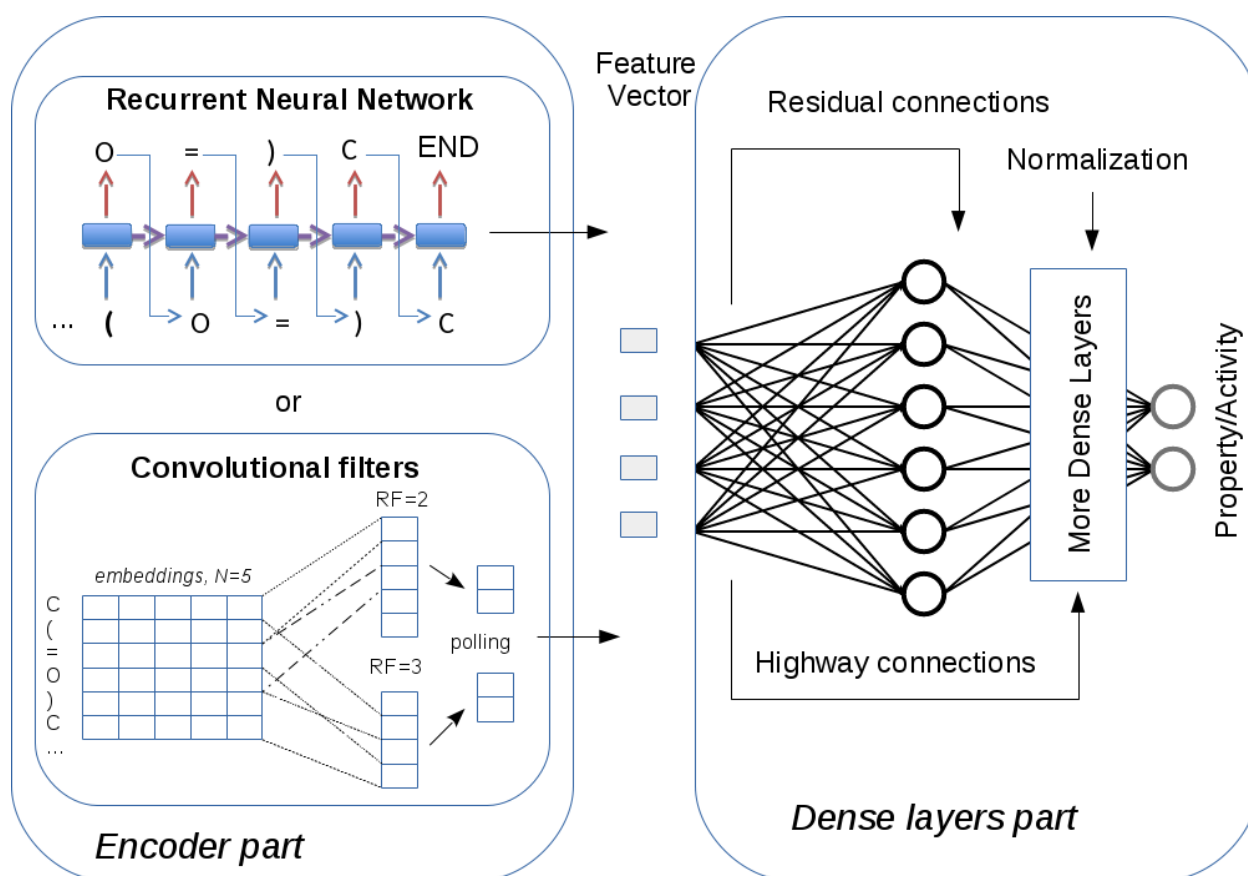


Fig. 2. Scheme of modern QSAR models based on ANN. The encoder part (left) extracts main features of the input data by means of RNN (top) or convolutional layers (bottom). Then the feature vector as usual

descriptors feeds to the dense layer part consisted of residual and highway connections, normalization layers, and dropouts.

resources to train the encoder part of a network. The concept of embeddings mitigates the problem by using the pre-trained weights designed for image [23] or text processing [24] tasks. It allows transfer learning from previous data and speeding up the training process for building models with significantly smaller datasets inaccessible for training from scratch. Typically, QSAR datasets contain only several hundreds of molecules, and SMILES-embeddings could improve models by developing better features.

One way of separately obtaining SMILES embeddings is to use classical autoencoder[25] approach where the input is the same as the output. In the case of SMILES, however, it would be more desirable to explore a variety of SMILES belonging to the same molecule due to redundant SMILES grammar, Fig. 1. We hypothesized that it is possible to train a neural network to conduct a SMILES canonization task in a Sequence-to-Sequence (Seq2Seq) manner like machine translation problem, where on the left side are non-canonical SMILES, and on the right side are their canonical equivalents. Recently, Seq2Seq was successfully applied to translation from InChi [26] codes to SMILES (Inchi2Sml) as well as from SMILES arbitrary to canonical SMILES (Sml2canSml), and to build QSAR models on extracted latent variables[27].

The state-of-the-art neural architecture for machine translation consists of stacked Long Short-Term Memory (LSTM) cells[28]. Training process for such networks has inherent for all kinds of Recurrent Neural Networks difficulties, i.g., vanishing gradients, and the impossibility of parallelization. Recently, a Transformer model [1] was proposed where all recurrent units are replaced with convolutional and element-wise feed-forward layers. The whole architecture shows a significant speed-up during training and inference with improved accuracy over translation benchmarks. The Transformer model was applied for prediction of reaction outcomes[29] and for retrosynthesis [30].

Our contributions in the article are as follows:

- we present a concept of dynamic SMILES embeddings that may be useful for a wide range of cheminformatics tasks;
- we scrutinize CharNN models based on these embeddings for regression and classification tasks and show that the method outperforms the state-of-the-art models;
- our implementation as well as source codes and SMILES-embeddings are available on

<https://github.com/bigchem/transformer-cnn>. We also provide ready-to-use implementation on <https://ochem.eu> within the OCHEM[3] environment.

Methods

SMILES canonization model

Dataset

To train the ANN to perform SMILES canonization task, we used the ChEMBL database[31] with length of SMILES less than or equal 110 characters (>93% of the entire database). The original dataset was augmented 10 times up to 17,657,995 canonization pairs written in reactions format separated by '>>'. Each pair contained on the left side a non-canonical, and on the right side – a canonical SMILES for the same molecule. Such an arrangement of the training dataset allowed us to re-use the previous Transformer code, which was originally applied for retrosynthesis task[30]. For completeness, we added for every compound a line where both left and right sides were identical, i.e. canonical SMILES, Fig. 3. Thus each molecule was present in the training set 11 times.




Fig. 3. Example of the data in the training file for canonization model of a small molecule CHEMBL351484. Every line contains a pair of non-canonical (left) and canonical (right) separated by ">>". One line has identical SMILES on both sides, stressed with the red box.

Model input

Seq2Seq models use one-hot encoding vector for the input. Its values are zero everywhere except the position of the current token which is set to one. Many works on SMILES use tokenization procedure[32, 33] that combines some characters, for example 'B' and 'r' to one token 'Br'. Other rules for handling most common two-letters elements, charges, and

stereochemistry also are used for preparing the input for the neural network. According to our experience, the use of more complicated schemes instead of simple character-level tokenization did not increase the accuracy of models[30]. Therefore a simple character-level tokenization was used in this study. The vocabulary of our model consisted of all possible characters from ChEMBL dataset and has 66 symbols:

`^#%()+-./0123456789=@ABCDEFGHIJKLMNPRSTVXYZ[\]abcdefghijklmnopqrstuvwxyz$`

Thus, the model could handle all diversity of drug-like compounds including stereochemistry, different charges, and inorganic ions. Two special characters were added to the vocabulary: ‘^’ to indicate the start of the sequence, and ‘\$’ to inform the model about the end of the data input.

Transformer model

The canonization model used in this work was based upon Transformer architecture consisting of two separate stacks of layers for the encoder and the decoder, respectively. Each layer incorporated some portion of knowledge written in its internal memory (V) with indexed access by keys (K). When new data arrived (Q), the layer calculated attention, to what it has already learned, and modified the input accordingly (see the original work on Transformers[1]), thus, forming the output of the self-attention layer and backlighting those parts that carry the essential information. Besides self-attention mechanism, the layer also contained several position-wise dense layers, normalization layer, and residual connections[1, 34]. Our model utilized three layers architecture of Transformer with 10 blocks of self-attention, i.e. the same one as used in our previous study[30]. After the encoding process was finished, the output of the top encoder layer contained a representation of a molecule suitable for decoding into canonical SMILES. In this study we used this representation as a well-prepared latent representation for QSAR modeling.

Tensorflow v1.12.02[35] was used as machine-learning framework to develop all parts of the Transformer, whereas RDKit v.2018.09.2[36] was used for SMILES canonization, and OpenBabel v2.3.1[15] for data augmentation.

QSAR model

We call the output of the Transformer's encoder part as a dynamic SMILES-embedding, Fig. 4. For a molecule with N -characters, the encoder produces the matrix with dimensions (N , $EMBEDDINGS$). Though technically this matrix is not an embedding because equivalent

characters have different values depending on position and surroundings, it can be considered so due to its role: to convert input one-hot raw vectors to real-value vectors in some latent space. Because these embeddings has variable lengths, we used a series of 1D convolutional filters as implemented in DeepChem[37] TextCNN method (<https://github.com/deepchem>).

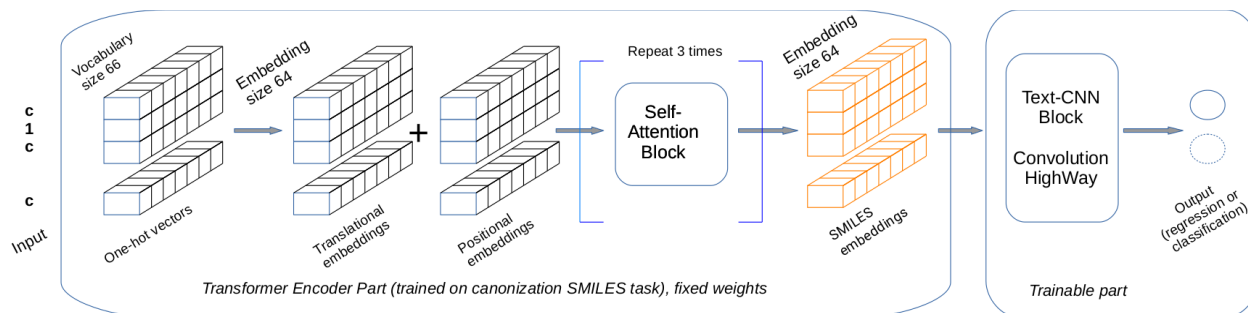


Fig. 4. The architecture of Transformer-CNN network.

Each convolution had a kernel size from the list [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 15, 20] and produced the following number of filters [100, 200, 200, 200, 200, 100, 100, 100, 100, 100, 160, 160], respectively. After GlobalMaxPool operation and the following concatenation of the pooling results, the data went through Dropout [38](rate=0.25), Dense(N=512), Highway[39] layers, and, finally, converted to the output layer which consisted of only one neuron for regression and two neurons for classification tasks. The weights of the Transformer's part were frozen in all experiments. All models used Adam optimizer with Mean Squared Error or Binary Cross-Entropy loss depending on the problem at hand. A fixed learning rate $\lambda = 10^{-4}$ was used. Early-stopping was used to prevent overfitting, select a best model and reduce training time. OCHEM calculations were performed using canonical SMILES as well as ten-times augmentation of SMILES during both training and prognosis. This number of SMILES augmentations was found to be an optimal one in our previous study[40]. An average value of the individual predictions for different representation of the same molecule were used as final model predictions to calculate statistical parameters.

The same five-fold cross-validation procedure was used to compare the models with results of our previous study[40]. The coefficients of determination[41]

$$r^2 = 1 - SS_{res}/SS_{tot} \quad (1)$$

where SS_{tot} is total variance of data and SS_{res} is residual unexplained variance of data was used to compare regression models and Area Under the Curve (AUC) was used for classification tasks.

Validation datasets

We used the same datasets (9 for regression and 9 for classification) that were exploited in our previous studies [11, 22]. Short information about these sets as well as links to original works are provided in Table 1. The datasets are available at OCHEM environment on <https://ochem.eu>.

Table 1. Descriptions of datasets used in the work.

Code	Description	Size	Code	Description	Size
Regression tasks			Classification tasks		
MP	Melting point [42]	19,104	HIV	Inhibition of HIV replication [43]	41,127
BP	Boiling point [44]	11,893	AMES	Mutagenicity [45]	6,542
BCF	Bioconcentration factor[44]	378	BACE	Human β -secretase 1 (BACE-1) inhibitors[43]	1,513
FreeSolv	Free solvation energy [43]	642	Clintox	Clinical trial toxicity [43]	1,478
LogS	Solubility [46]	1,311	Tox21	In-vitro toxicity [43]	7,831
Lipo	Lipophilicity [47]	4,200	BBBP	Blood-brain barrier [43]	2,039
BACE	IC50 of human β -secretase 1 (BACE-1) inhibitors [43]	1,513	JAK3	Janus kinase 3 inhibitor[48]	886
DHFR	Dihydrofolate reductase inhibition[49]	739	BioDeg	Biodegradability [50]	1,737
LEL	Lowest effect level [51]	483	RP AR	Endocrine disruptors [52]	930

Results and discussion

SMILES canonization model

The Transformer model was trained for 10 epochs with learning rate changing according to the formula:

$$\lambda = factor * \min(1.0, step / warmup) / \max(step, warmup)(1)$$

where $factor = 20$, $warmup = 16,000$ steps, and if $\lambda < 10^{-4}$ then $\lambda = 10^{-4}$. The settings for the learning rate were similar to those used in our retro-synthesis study. Each epoch contained

275,907 steps (batches). No early-stopping or weights averaging was applied. Learning curves are shown in Fig. 5.

To validate the model, we sampled 500,000 ChEMBL-like SMILES (only 8,617 (1.7%) of them were canonical) from a generator[53] and checked how accurately the model can restore canonical SMILES for these molecules. We intentionally selected the generated SMILES keeping in mind possible application of the proposed method in the artificial intelligence-driven pipelines of *de-novo* development of new drugs. The model correctly canonized **83,6%** of all samples, Table 2.

Table 2. Validation of canonization model.

Strings	All	Correctly canonized
All	500,000	418,233 (83.6%)
Stereo (with @)	77,472	28,821 (37,2%)
Cis/trans (with / or \)	54,727	40,483 (73,9%)




Fig. 5. Learning curves: 1) learning rate schedule (axes bottom and right), and 2) character-based accuracy (axes bottom and left) on the training dataset for the first four epochs.

QSAR modeling

For the QSAR modelling the saved embedding was used. The training was done using fixed learning rate $\lambda = 0.001$ for $n=100$ epochs. The early stopping with 10% of randomly selected SMILES was used to identify the optimal model. Table 2, Fig. 6 compare results for regression datasets while Table 3, Fig. 7 does it for classification tasks. The standard mean errors of the values were calculated using bootstrap procedure as explained elsewhere [50].

With an exception of few datasets, the proposed method provided similar or better results than those of calculated using descriptor-based approaches as well as of the other SMILES-based approaches investigated in our previous study.[40] The data augmentation was critically important for the Transformer-CNN method to achieve its high performance. We used augmentation $n=10$, i.e., 10 SMILES were randomly generated and used for model development and application, which was found as an optimal one in the aforementioned previous study.

Table 2. Coefficient of determination, r^2 , calculated for regression sets (higher values are better)¹

Dataset	Descriptor based methods ²	SMILES based (augm=10) ²	Transformer-CNN, no augm.	Transformer-CNN, augm=10	CDDD descriptors ³
MP	0.83	0.85	0.83	0.86	0.85
BP	0.98	0.98	0.97	0.98	0.98
BCF	0.85	0.85	0.71±0.02	0.85	0.81
FreeSolv	0.94	0.93	0.72±0.02	0.91	0.93
LogS	0.92	0.92	0.85	0.91	0.91
Lipo	0.7	0.72	0.6	0.73	0.74
BACE	0.73	0.72	0.66	0.76	0.75
DHFR	0.62±0.03	0.63±0.03	0.46±0.03	0.67±0.03	0.61±0.03
LEL	0.19±0.04	0.25±0.03	0.2±0.03	0.27±0.04	0.23±0.04

1-We omitted the standard mean errors, which are 0.01 or less, for the reported values. 2- results from our previous study [22]. 3 - Best performance calculated with CDDD descriptors obtained using autoencoder Sml2canSml from [27].




Fig. 6. Coefficient of determination, r^2 , calculated for regression sets (higher values are better).

Table 3. AUC calculated for classification sets (higher values are better)

Dataset	Descriptor based methods ²	SMILES based (augm=10) ²	Transformer-CNN, no augm.	Transformer-CNN, augm=10	CDDD descriptors ³
HIV	0.82	0.78	0.81	0.83	0.74
AMES	0.86	0.88	0.86	0.89	0.86
BACE	0.88	0.89	0.89	0.91	0.9
Clintox	0.77±0.03	0.76±0.03	0.71±0.02	0.77±0.02	0.73±0.02
Tox21	0.79	0.83	0.81	0.82	0.82
BBBP	0.90	0.91	0.9	0.92	0.89
JAK3	0.79±0.02	0.8±0.02	0.70±0.02	0.78±0.02	0.76±0.02
BioDeg	0.92	0.93	0.91	0.93	0.92
RP AR	0.85	0.87	0.83	0.87	0.86

1-We omitted the standard mean errors, which are 0.01 or less, for the reported values. 2- results from our previous study [22] . 3 - Best performance calculated with CDDD descriptors obtained using Sml2canSml autoencoder from [27].




Fig.7. AUC calculated for classification sets (higher values are better).

Similar to Transformer-CNN the Sml2canSml used internal representation, which was developed from mapping of arbitrary SMILES to canonical SMILES. The difference was that Sml2canSml generated a fixed set of 512 latent variables (CDDD descriptors), while Transformer-CNN representation had about the same length as the initial SMILES. Sml2canSml CDDD could be used as descriptors with any traditional machine learning methods while Transformer-CNN required convolutional neural network to process variable length output and to correlate it with the analysed properties. Sml2canSml was added as CDDD descriptors to OCHEM. These descriptors were analysed by the same methods as used in the previous work, i.e., LibSVM[54], Random Forest[55], XGBoost[56] as well as by Associative Neural Networks (ASNN)[57] and Deep Neural Networks[58]. Exactly the same protocol, 5 fold cross-validation, was used for all calculations. The best performance using the CDDD descriptors was calculated by ASNN and LibSVM methods, which contributed models with the highest accuracies for seven and five datasets, respectively (LibSVM method provided the best performance in the original study). Transformer-CNN provided better or similar results compared to the CDDD descriptors for all datasets with an exception of Lipo and FreeSolv. It should be also mentioned, that CDDD descriptors could only process molecules which satisfy the following conditions:

$$\log P \in (-5,7) \text{ and}$$

mol_weight \in (12,600) and
num_heavy_atoms \in (3, 50) and
molecule is organic.

These limitations appeared due to the preparation of the training set to develop Sml2canSml encoder. The limitations resulted in the exclusion of a number of molecules, which failed one or several of the above conditions. Contrary to Sml2canSml encoder, we trained Transformer-CNN with very diverse molecules from ChEMBL and thus the developed models could be applied to any molecule, which is processed by RDKit. Actually, exclusion of molecules, for which CDDD descriptors failed to be calculated, did not significantly change results of Transformer models: some models improved while others decreased their accuracy for about ~ 0.01 respective performance values. For example, for Lipo and FreeSolv sets the accuracy of Transformer-CNN model increased to $r^2 = 0.92$ and 0.75 , respectively while for BBB AUC decreased to 0.91 .

Conclusions and outlook

We propose a SMILES canonization model based on Transformer architecture that extracts information-rich real-value embeddings during the encoding process and exposes them for further QSAR-oriented blocks to model biological activity or physico-chemical properties. TextCNN approaches can efficiently work with these embeddings, and the final quality of the QSAR models is higher compared to the models obtained with the state-of-the-art methods on the majority of diverse benchmark datasets. The Transformer-CNN architecture requires less than a hundred iterations to converge for new tasks. It can be easily embedded into *de-novo* drug development pipelines. The code is available on <https://github.com/bigchem/transformer-cnn> as well as on-line version on <http://ochem.eu>.

The method developed predicts the endpoint based on an average of individual prognosis for a batch of augmented SMILES belonging to the same molecule. The deviation within the batch can serve as a measure of a confidence interval of the prognosis, whereas the possibility to canonize SMILES can be used for deriving applicability domains of models. These questions will be addressed in the upcoming studies.

Abbreviations

ANN: Artificial Neural Network; CNN: Convolutional Neural Network; LSTM: Long Short-Term memory; OCHEM: On-line chemical database and modeling environment; SMILES: Simplified

Molecular-Input Line-Entry System; QSAR/QSPR: Quantitative Structure ActivityProperty Relationship; RF: Receptive Field; RNN: Recurrent Neural Network.CNN: Convolutional Neural Network

Declarations

Availability of data and materials

The source code of Transformer-CNN is available on <https://github.com/bigchem/transformer-cnn>. Ready-to-use implementation as well as training datasets, and final QSAR models are stored on <https://ochem.eu> within the OCHEM environment.

Competing interests

The authors declare that they have no actual or potential conflicts of interests.

Funding

This study was funded by the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 676434, "Big Data in Chemistry" and ERA-CVD "CardioOncology" project, BMBF 01KL1710.

Authors' contributions

PK implemented the method, IVT and GC performed the analysis and benchmarking. All authors interpreted results, read and approved the manuscript.

Acknowledgments

The authors thank NVIDIA Corporation for donating Quadro P6000, Titan Xp, and Titan V graphics cards for this research work.

References

1. Vaswani A, Shazeer N, Parmar N, et al (2017) Attention Is All You Need. ArXiv
2. Zhang X, Zhao J, LeCun Y (2015) Character-level Convolutional Networks for Text Classification. arXiv [cs.LG]

3. [Sushko I, Novotarskyi S, Körner R, et al \(2011\) Online chemical modeling environment \(OCHEM\): web platform for data storage, model development and publishing of chemical information. J Comput Aided Mol Des 25:533–554.
https://doi.org/10.1007/s10822-011-9440-2](#)
4. Mauri A, Consonni V, Pavan M, Todeschini R (2006) Dragon software: An easy approach to molecular descriptor calculations. Match 56:237–248
5. Baskin I, Varnek A (2009) Fragment descriptors in SAR/QSAR/QSPR studies, molecular similarity analysis and in virtual screening. ChemInform 40:i
6. [Eklund M, Norinder U, Boyer S, Carlsson L \(2014\) Choosing feature selection and learning algorithms in QSAR. J Chem Inf Model 54:837–843.
https://doi.org/10.1021/ci400573c](#)
7. [Chen H, Engkvist O, Wang Y, et al \(2018\) The rise of deep learning in drug discovery. Drug Discov Today 23:1241–1250.
https://doi.org/10.1016/j.drudis.2018.01.039](#)
8. Duvenaud D, Maclaurin D, Aguilera-Iparraguirre J, et al (2015) Convolutional Networks on Graphs for Learning Molecular Fingerprints. arXiv [cs.LG]
9. [Coley CW, Barzilay R, Green WH, et al \(2017\) Convolutional Embedding of Attributed Molecular Graphs for Physical Property Prediction. J Chem Inf Model 57:1757–1772.
https://doi.org/10.1021/acs.jcim.6b00601](#)
10. [Gómez-Bombarelli R, Wei JN, Duvenaud D, et al \(2018\) Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. ACS Central Science 4:268–276.
https://doi.org/10.1021/acscentsci.7b00572](#)
11. Kimber TB, Engelke S, Tetko IV, et al (2018) Synergy Effect between Convolutional Neural Networks and the Multiplicity of SMILES for Improvement of Molecular Prediction. arXiv
12. Gilmer J, Schoenholz SS, Riley PF, et al (2017) Neural Message Passing for Quantum Chemistry. arXiv [cs.LG]
13. Shang C, Liu Q, Chen K-S, et al (2018) Edge Attention-based Multi-Relational Graph Convolutional Networks. arXiv [stat.ML]
14. [Weininger D \(1988\) SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. J Chem Inf Comput Sci 28:31–36.
https://doi.org/10.1021/ci00057a005](#)
15. [O'Boyle NM, Banck M, James CA, et al \(2011\) Open Babel: An open chemical toolbox. J Cheminform 3:33.
https://doi.org/10.1186/1758-2946-3-33](#)
16. [Vidal D, Thormann M, Pons M \(2005\) LINGO, an efficient holographic text based method to calculate biophysical properties and intermolecular similarities. J Chem Inf Model 45:386–393.
https://doi.org/10.1021/ci0496797](#)
17. Zhang X, LeCun Y (2015) Text Understanding from Scratch. arXiv [cs.LG]
18. Goh GB, Hodas NO, Siegel C, Vishnu A (2017) SMILES2Vec: An Interpretable General-Purpose Deep Neural Network for Predicting Chemical Properties. arXiv [stat.ML]

19. Jastrzębski S, Leśniak D, Czarnecki WM (2016) Learning to SMILE(S). arXiv [cs.CL]
20. Goh GB, Siegel C, Vishnu A, Hodas NO (2017) Using Rule-Based Labels for Weak Supervised Learning: A ChemNet for Transferable Chemical Property Prediction. arXiv [stat.ML]
21. [Zheng S, Yan X, Yang Y, Xu J \(2019\) Identifying Structure-Property Relationships through SMILES Syntax Analysis with Self-Attention Mechanism. J Chem Inf Model 59:914–923.
https://doi.org/10.1021/acs.jcim.8b00803](#)
22. Igor V. Tetko, Pavel Karpov, Eric Bruno, Talia B. Kimber, Guillaume Godin Augmentation is What You Need! In: Igor V. Tetko, Fabian Theis, Pavel Karpov, Vera Kurkova (ed) 28th International Conference on Artificial Neural Networks, Munich, Germany, September 17–19, 2019 Proceedings, Part V, Workshop and Special sessions. Springer
23. Kiela D, Bottou L (2014) Learning image embeddings using convolutional neural networks for improved multi-modal semantics. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP). pp 36–45
24. Pennington J, Socher R, Manning CD (2014) Glove: Global vectors for word representation. In: In EMNLP
25. [Hinton GE, Salakhutdinov RR \(2006\) Reducing the dimensionality of data with neural networks. Science 313:504–507.
https://doi.org/10.1126/science.1127647](#)
26. [Heller S, McNaught A, Stein S, et al \(2013\) InChI - the worldwide chemical structure identifier standard. J Cheminform 5:7.
https://doi.org/10.1186/1758-2946-5-7](#)
27. [Winter R, Montanari F, Noé F, Clevert D-A \(2019\) Learning continuous and data-driven molecular descriptors by translating equivalent chemical representations. Chem Sci 10:1692–1701.
https://doi.org/10.1039/c8sc04175j](#)
28. [Hochreiter S, Schmidhuber J \(1997\) Long Short-Term Memory. Neural Comput 9:1735–1780.
https://doi.org/10.1162/neco.1997.9.8.1735](#)
29. Schwaller P, Laino T, Gaudin T, et al (2018) Molecular Transformer for Chemical Reaction Prediction and Uncertainty Estimation. arXiv
30. Pavel Karpov, Guillaume Godin, Igor V. Tetko A Transformer Model for Retrosynthesis. In: Igor V. Tetko, Fabian Theis, Pavel Karpov, Vera Kurkova (ed) 28th International Conference on Artificial Neural Networks, Munich, Germany, September 17–19, 2019 Proceedings, Part V, Workshop and Special sessions. Springer
31. [Gaulton A, Bellis LJ, Bento AP, et al \(2012\) ChEMBL: a large-scale bioactivity database for drug discovery. Nucleic Acids Res 40:D1100–7.
https://doi.org/10.1093/nar/gkr777](#)
32. Segler MHS, Kogej T, Tyrchan C, Waller MP (2017) Generating Focussed Molecule Libraries for Drug Discovery with Recurrent Neural Networks
33. Gupta A, Müller AT, Huisma BJH, et al (2018) Generative Recurrent Networks for De Novo Drug Design. Mol Inform 37:1700111
34. Rush A (2018) The Annotated Transformer. In: Proceedings of Workshop for NLP Open

Source Software (NLP-OSS). pp 52–60

35. Abadi M, Barham P, Chen J, et al (2016) TensorFlow: A system for large-scale machine learning
36. [Landrum G RDKit: Open-source cheminformatics. http://www.rdkit.org](http://www.rdkit.org)
37. Ramsundar B, Eastman P, Walters P, Pande V (2019) Deep Learning for the Life Sciences: Applying Deep Learning to Genomics, Microscopy, Drug Discovery, and More. "O'Reilly Media, Inc."
38. Srivastava N, Hinton G, Krizhevsky A, et al (2014) Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *J Mach Learn Res* 15:1929–1958
39. Srivastava RK, Greff K, Schmidhuber J (2015) Highway Networks. *arXiv [cs.LG]*
40. Tetko IV, Karpov P, Bruno E, et al (2019) Augmentation Is What You Need!: 28th International Conference on Artificial Neural Networks, Munich, Germany, September 17–19, 2019, Proceedings. In: Tetko IV, Kůrková V, Karpov P, Theis F (eds) *Artificial Neural Networks and Machine Learning – ICANN 2019: Workshop and Special Sessions*. Springer International Publishing, Cham, pp 831–835
41. Draper NR, Smith H (2014) *Applied Regression Analysis*. John Wiley & Sons
42. [Tetko IV, Sushko Y, Novotarskyi S, et al \(2014\) How accurately can we predict the melting points of drug-like compounds? *J Chem Inf Model* 54:3320–3329. https://doi.org/10.1021/ci5005288](https://doi.org/10.1021/ci5005288)
43. Wu Z, Ramsundar B, Feinberg EN, et al (2018) MoleculeNet: a benchmark for molecular machine learning. *Chemical Science* 9:513–530
44. Brandmaier S, Sahlin U, Tetko IV, Öberg T (2012) PLS-Optimal: A Stepwise D-Optimal Design Based on Latent Variables. *Journal of Chemical Information and Modeling* 52:975–983
45. Sushko I, Novotarskyi S, Körner R, et al (2010) Applicability Domains for Classification Problems: Benchmarking of Distance to Models for Ames Mutagenicity Set. *Journal of Chemical Information and Modeling* 50:2094–2111
46. [Delaney JS \(2004\) ESOL: estimating aqueous solubility directly from molecular structure. *J Chem Inf Comput Sci* 44:1000–1005. https://doi.org/10.1021/ci034243x](https://doi.org/10.1021/ci034243x)
47. Huuskonen JJ, Livingstone DJ, Tetko IV IV (2000) Neural network modeling for estimation of partition coefficient based on atom-type electrotopological state indices. *J Chem Inf Comput Sci* 40:947–955
48. Suzuki K, Nakajima H, Saito Y, et al (2000) Janus kinase 3 (Jak3) is essential for common cytokine receptor γ chain (γ c)-dependent signaling: comparative analysis of γ c, Jak3, and γ c and Jak3 double-deficient mice. *International Immunology* 12:123–132
49. Sutherland JJ, Weaver DF (2004) Three-dimensional quantitative structure-activity and structure-selectivity relationships of dihydrofolate reductase inhibitors. *Journal of Computer-Aided Molecular Design* 18:309–331

50. [Vorberg S, Tetko IV \(2014\) Modeling the Biodegradability of Chemical Compounds Using the Online CHEmical Modeling Environment \(OCHEM\). Mol Inform 33:73–85.
https://doi.org/10.1002/minf.201300030](#)
51. [Novotarskyi S, Abdelaziz A, Sushko Y, et al \(2016\) ToxCast EPA in Vitro to in Vivo Challenge: Insight into the Rank-I Model. Chem Res Toxicol 29:768–775.
https://doi.org/10.1021/acs.chemrestox.5b00481](#)
52. Rybacka A, Rudén C, Tetko IV, Andersson PL (2015) Identifying potential endocrine disruptors among industrial chemicals and their metabolites – development and evaluation of in silico tools. Chemosphere 139:372–378
53. Zhonghua Xia, Pavel Karpov, Grzegorz Popowicz, Igor V. Tetko (2019) Focused Library Generator: Case of Mdmx Inhibitors. Accepted to print
54. [Chang C-C, Lin C-J \(2011\) LIBSVM: A library for support vector machines. ACM Transactions on Intelligent Systems and Technology \(TIST\) 2:27.
https://doi.org/10.1145/1961189.1961199](#)
55. [Breiman L \(2001\) Random Forests. Mach Learn 45:5–32.
https://doi.org/10.1023/A:1010933404324](#)
56. Chen T, Guestrin C (2016) XGBoost: A Scalable Tree Boosting System. arXiv [cs.LG]
57. [Tetko IV \(2002\) Associative Neural Network. Neural Process Letters 16:187–199.
https://doi.org/10.1023/A:1019903710291](#)
58. [Sosnin S, Karlov D, Tetko IV, Fedorov MV \(2019\) Comparative Study of Multitask Toxicity Modeling on a Broad Chemical Space. J Chem Inf Model 59:1062–1072.
https://doi.org/10.1021/acs.jcim.8b00685](#)

Trans-CNN.docx (478.64 KiB)

[view on ChemRxiv](#) • [download file](#)
