

# SMILES2vec

An Interpretable General-Purpose Deep Neural Network for Predicting Chemical Properties

Garrett B. Goh

PNNL

garrett.goh@pnnl.gov

Charles Siegel

PNNL

charles.siegel@pnnl.gov

Nathan Hodas

PNNL

nathan.hodas@pnnl.gov

Abhinav Vishnu

PNNL

abhinav.vishnu@pnnl.gov

## ABSTRACT

Chemical databases store information in text representations, and the SMILES format is a universal standard used in many cheminformatics software. Encoded in each SMILES string is structural information that can be used to predict complex chemical properties. In this work, we develop SMILES2vec, a deep RNN that automatically learns features from SMILES to predict chemical properties, without the need for additional explicit feature engineering. Using Bayesian optimization methods to tune the network architecture, we show that an optimized SMILES2vec model can serve as a general-purpose neural network for predicting distinct chemical properties including toxicity, activity, solubility and solvation energy, while also outperforming contemporary MLP neural networks that uses engineered features. Furthermore, we demonstrate proof-of-concept of interpretability by developing an explanation mask that localizes on the most important characters used in making a prediction. When tested on the solubility dataset, it identified specific parts of a chemical that is consistent with established first-principles knowledge with an accuracy of 88%. Our work demonstrates that neural networks can learn technically accurate chemical concept and provide state-of-the-art accuracy, making interpretable deep neural networks a useful tool of relevance to the chemical industry.

## CCS CONCEPTS

•**Computing methodologies** → **Natural language processing; Neural networks**; Transfer learning; •**Applied computing** → **Chemistry**; *Computational biology*; *Bioinformatics*; •**Theory of computation** → *Bayesian analysis*;

## KEYWORDS

Natural Language Processing, Bayesian Optimization, Interpretability, Transfer Learning, Cheminformatics

### ACM Reference format:

Garrett B. Goh, Nathan Hodas, Charles Siegel, and Abhinav Vishnu. 2018. SMILES2vec. In *Proceedings of ACM SIGKDD Conference, London, UK, Aug 2018 (KDD 2018)*, 8 pages.

DOI: 10.475/123.4

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

KDD 2018, London, UK

© 2018 Copyright held by the owner/author(s). 123-4567-24-567/08/06...\$15.00  
DOI: 10.475/123.4

## 1 INTRODUCTION

In the chemical industry, designing chemicals with desired characteristics is a bottleneck in the development of new products. Despite decades of research, much of modern day chemical design is still driven by serendipity and chemical intuition, although improvements to rational chemical design has been incrementally improving over time. [23] Currently, there exists two key contributing factors in the current state in physics-based or rule-based computational chemistry methods that accounts for the gap towards true rational chemical design. One factor is driven by technical limitations, such as that in compute capacity, and the other is incomplete or partial understanding of the underlying chemical concepts. The first factor is being addressed with development of GPU-accelerated code for molecular modeling [36] and special-purpose chemistry supercomputers [34], but the second factor requires solutions grounded in further fundamental research, which is often a slower process.

### 1.1 Limitations of Feature Engineering and Black Box Models

For the prediction of chemical properties that cannot be easily computed through physics-based or rule-based methods, modern *in-silico* modeling in chemistry is therefore predicated on correlating engineered features with the activity or property of the chemical, which is formally known as the field of Quantitative Structure-Activity or Structure-Property Relationship (QSAR/QSPR) modeling [4]. Feature engineering in chemistry is a sophisticated science that stretches back to the late 1940s [30]. Molecular descriptors, as they are termed by chemists, are basic computable properties or sophisticated descriptions of a chemical's structure, and these engineered features were developed based on first-principles knowledge. To date, over 5000 molecular descriptors have been developed [38]. In addition, molecular fingerprints have also been designed, which instead of computing a basic property, provides a description of a specific part of the chemical's structure [33].

Since the 1980s, various machine learning (ML) algorithms have been used to predict the activity or property of chemicals [4] using molecular descriptors and/or fingerprints as input features. More recently, deep learning (DL) models have also been developed [9, 21, 28, 31]. In general, these models either perform at parity or slightly outperform prior state-of-the-art models based on traditional ML algorithms for chemical applications [14].

Compared to computer vision (CV) and natural language processing research (NLP), the use of DL models in chemistry relies heavily on these engineered features. This may be problematic as it limits the neural network’s search space of potentially learnable representations. This is further exacerbated in situations in which engineered features are not appropriate or inadequate due to the lack of well-developed domain knowledge, which originates from the second factor that limits the impact of existing computational chemistry methods. In contrast, the dominant approach in CV/NLP research is to train DL models directly on typically unaltered raw data of large datasets with little or minimal feature engineering research. For example, unaltered images are used as the input in various CNN models [18] and similarly unaltered text is used in LSTM-based models [41]. Therefore, developing DL models that leverage on representation learning is a logical advance for the field of chemistry as well.

Lastly, an additional challenge associated with current ML/DL models is the lack of interpretability. Typically operated as opaque black-box models, it is difficult to gain any scientific understanding as to *why or how* an algorithm predicts a particular chemical property. For typical applications in CV/NLP research, this may not be an issue. However, for the chemical industry, particularly for regulated products, such as requesting FDA approval for new drugs, an explanation of how the chemical works (i.e. interacts with the body is a requirement). Therefore, models that have increased interpretability or explainability is not only of industrial relevance, but it will also enable chemists to formulate new hypothesis to improve on and possibly accelerate the pace fundamental research.

## 1.2 Related Work

In chemistry, raw unaltered data would typically refer to a representation that describes the structure and orientation of a chemical. In basic chemistry education, students are taught how to draw a 2D diagram of a chemical (i.e. an image), which also serves as the primary medium of communication amongst chemist. Alternatively, the same structural information can be encoded as graphs. Indeed, convolutional neural network (CNN) models that use chemical images [16, 17, 39] and other DL models that use molecular graphs [12, 24] have been recently developed. In addition, a chemical’s structural information can also be encoded in text format, such as SMILES [40], which is also the basis for interoperability between various cheminformatics software packages. In terms of text representations, we acknowledge there has been some prior work in this direction [2, 22]. In terms of interpretable DL models, while we have seen advances in conventional CV/NLP applications [32], at the time of writing, we are not aware of any interpretable DL models in chemistry that learns directly on raw data.

## 1.3 Contributions

Our work improves the existing state of learning directly from chemical text representations, and it is also the first interpretable neural network that works on chemical text. In the process, our work also addresses the following question: *Is the SMILES representation sufficient to capture the first order distinction between different chemical properties?* Assuming that the above hypothesis is true, *would it then be possible to validate what a neural network learns with*

*established first principles knowledge on simple chemical properties, such as solubility?* Our contributions are as follows:

- We perform extensive experiments to determine the optimal neural network architecture for interpreting the SMILES “chemical language”.
- We developed an explanation mask to explain *why or how* the neural network makes a particular prediction.
- We show how an optimal SMILES2vec network architecture can be generalized to predict broad range of properties that are of relevance to multiple industries, including pharmaceuticals, biotechnology, materials and consumer goods.
- We demonstrate that SMILES2vec models, despite having no feature engineering, achieves better accuracy than contemporary multi-layer perceptron (MLP) models that uses engineered features.

The organization for the rest of the paper is as follows. In section 2, we examine the datasets, its broad applicability to chemical-affiliated industries, as well as the Bayesian optimization process used for refining the network architecture, and the training protocols of the neural network. Then, in section 3, we document the experiments used to develop the final SMILES2vec network architecture and constructing of the interpretability masks. Lastly, in section 4, we perform additional experiments to quantify the accuracy of interpretability using the solubility dataset as an example, and evaluate SMILES2vec accuracy of chemical property predictions against contemporary DL models.

## 2 METHODS

Here, we document the methods used in the development of SMILES2vec. First, we provide a brief introduction to SMILES. Then, we provide details on the datasets used, data splitting and preparation. Then, we examine the details of refining the neural network architecture using Bayesian methods, as well as the training protocol and evaluation metrics for the neural network.

### 2.1 Introduction to SMILES

SMILES is a “chemical language” [40] that encodes structural information of a chemical into a compact text representation. There is a regular grammar to SMILES. For example, the alphabets denote atoms, and in some cases also what type of atoms. For example, c and C denote aromatic and aliphatic carbons respectively. Special characters like ‘=’ denote the type of bonds (connections between atoms). Rings are denoted by encapsulating numbers, and side chains by round brackets. Thus, with sufficient training a chemist can read SMILES and infer the structure of the chemical. From this structural information, more complex properties can be predicted.

Inspired by language translation RNN work [41], we do not explicitly encode information about the grammar of SMILES. Instead, we expect that the RNN should learn these patterns and if necessary use them to develop intermediate features that would be relevant for predicting a variety of chemical properties.

### 2.2 Datasets Used

Our work creates a RNN model for general chemical property prediction, and ideally it should work effectively for different types of

Dataset	Property	Task	Size
Tox21	Non-Physical (Toxicity)	Multi-task classification	8014
HIV	Non-Physical (Activity)	Single-task classification	41,193
FreeSolv	Physical (Solvation)	Single-task regression	643
ESOL	Physical (Solubility)	Single-task regression	1128

**Table 1: Characteristics of the 4 datasets used to evaluate the performance of SMILES2vec.**

properties without significant network topology changes. To ensure that our results are comparable with contemporary DL models reported in the literature [42] and earlier work on Chemception CNN models [16, 17], we used the Tox21, HIV, and FreeSolv dataset from the MoleculeNet benchmark [42] for predicting toxicity, activity and solvation free energy respectively. These datasets (see Table 1) represent a good mix of dataset sizes, type of chemical properties and regression vs classification tasks. In addition, we also used the ESOL solubility dataset to evaluate the interpretability of SMILES2vec.

### 2.3 Relevance to Chemical Industries

In terms relevance to chemical-affiliated industries, toxicity prediction has importance for chemicals that require FDA approval, which includes drugs and other therapeutics (pharmaceuticals) as well as cosmetics (consumer goods). [27] Activity is a measurement of how well a chemical binds to its intended target and is one of the factors that determine how well a chemical may perform as a drug. [3] Therefore, accurate activity predictions are of relevance to both pharmaceuticals and biotechnology industries. Predicting solubility is an important consideration for developing formulations for products relevant to pharmaceuticals and consumer goods, [11] and it also affects the bioavailability of drugs. Lastly, free energy values itself are computable by physics-based simulations, and such methods are currently being employed by pharmaceuticals, consumer goods and materials industries. [7]

### 2.4 Data Preparation

The length of the SMILES string directly impacts the compute resources required to train RNN models. To maintain a balance between maximum amount of SMILES data, but also rapid training time, we surveyed the ChEMBL database, a collection of industrially-relevant chemicals that has over 1 million entries. [13] Using this database as a proxy for relevant chemicals, we calculated that setting a maximum length of 250 characters would encompass 99.9% of existing entries. Therefore, in the above-listed datasets, we excluded entries of more than 250 characters in the dataset.

Next, we created a dictionary that mapped the unique characters as one-hot encodings. Zero padding was also applied to ensure that shorter strings had a uniform size of 250 characters. In addition, extra padding of 10 zeroes were added both to the left and right of

the string. Apart from the above-mentioned steps, no additional data augmentation steps were performed.

### 2.5 Data Splitting

We used a dataset splitting approach that is similar to that reported in previous work [16]. A separate test set was first partitioned out to serve as a test for model generalizability. For the Tox21 and HIV dataset, 1/6th was partitioned out to form the test set, and for the FreeSolv and ESOL dataset, 1/10th was used to create the test set. The remaining 5/6th or 9/10th of the dataset was then used in the random 5-fold cross validation approach for training.

Model performance and early stopping criterion was determined by validation loss. Lastly, we oversampled the minority class for classification tasks (Tox21, HIV) to mitigate class imbalance. This was achieved by computing the ratio of both classes, and appending additional data from the smaller class by that ratio. The oversampling step was performed after stratification, to ensure that the same molecule is not repeated across training/validation/test sets.

### 2.6 Bayesian Optimization of Neural Network Design

We used a Bayesian optimizer, SigOpt [10] to optimize the hyperparameters related to the neural network topology. Each different set of network hyperparameters are defined as a separate trial, and for each trial, we trained the model to completion using a standardized supervised training protocol. After each trial, the validation metric (AUC for classification tasks, RMSE for regression tasks) was used as input to the Bayesian optimizer for suggesting new network designs.

To prevent overfitting during this optimization process, the splitting of the dataset between training and validation sets was governed by a random seed. However, a fixed test set was maintained throughout, and this is also not used in the optimization process. By comparing the difference in validation and test set metrics, it would thus allow us to determine if the network design was being overfitted to the training/validation data. No hyperparameters optimization was performed for the learning protocol. Lastly, it should be noted that only a subset of the dataset was used in the Bayesian optimization. Specifically, we used only a single task (nr-ahr toxicity) from the Tox21 dataset and the FreeSolv dataset.

### 2.7 Training the Neural Network

SMILES2vec was trained using a Tensorflow backend [1] with GPU acceleration using NVIDIA CuDL libraries. [5] The network was created and executed using the Keras 2.0 functional API interface [8]. The RMSprop algorithm [19] was used to train for 250 epochs using the standard settings recommended (learning rate =  $10^{-3}$ ,  $\rho = 0.9$ ,  $\epsilon = 10^{-8}$ ). The batch size was 32, and we also included early stopping to reduce overfitting. This was done by monitoring the loss of the validation set, and if there was no improvement in the validation loss after 25 epochs, the last best model was saved as the final model.

For classification tasks, we used the binary crossentropy loss function for training. The performance metric reported in our work is area under the ROC curve (AUC). For regression tasks, we used the mean average error as the loss function for training. The

performance metric reported is RMSE. Unless specified otherwise, the reported results in our work denote the mean value of the performance metric, obtained from the 5 runs in the 5-fold cross validation.

### 3 EXPERIMENTS

In this section, we first conduct several Bayesian optimization experiments to optimize SMILES2vec’s architecture and hyperparameters. Then, we conduct further experiments to develop an explanation mask for improving interpretability of the model.

#### 3.1 SMILES2vec Neural Network Design

RNNs, particularly those based on LSTMs [20] or GRUs [6] are effective neural network designs for learning from text data. Its effectiveness has been demonstrated in examples like the Google Neural Translation Machine that uses an architecture of 8+8 layers of residual LSTM units [41]. Most of the other reported application of RNNs in NLP research are similarly used to model sequence-to-sequence predictions, and often fewer (i.e. 2 to 4) layers have been found to be sufficiently accurate for their tasks.

Our work differs from conventional NLP research as we are modeling sequence-to-vector predictions, where the sequence is a SMILES string, and the vector is a measured chemical property. Because of this, and also because SMILES is a fundamentally different language, commonly-used techniques in NLP research, such as embeddings like Word2vec [29] cannot be easily adapted for use in our work. Therefore, a substantial component of our work is in the design of the RNN architecture specific to SMILES.

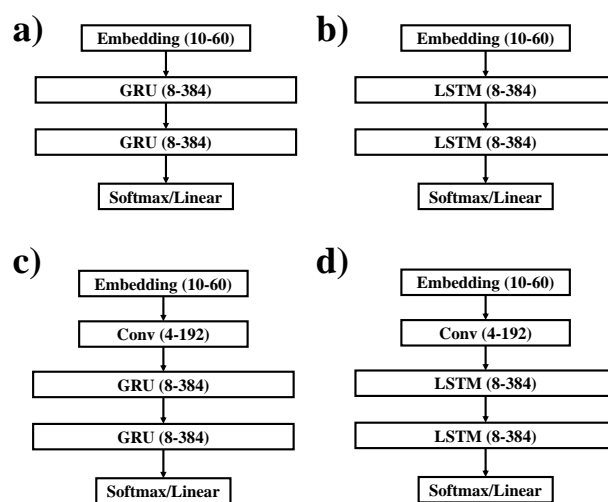
#### 3.2 Architectural Class Exploration

We first explore the RNN model’s architecture class, which primarily includes high-level design choices, such as the type of units used, type of layers, arrangement of layers, etc. LSTMs and GRUs are the two major RNN units used in the literature, and form the basis of two architectural classes. The template design for each class starts with an embedding layer that feeds into a 2-layer bidirectional GRU or 2-layer bidirectional LSTM as illustrated in Figure 1. In addition, we explored the utility of adding a 1D convolutional layer between the embedding and GRU/LSTM layers. This design forms the template of the other two architectural classes explored.

A separate Bayesian optimization was used to optimize the hyperparameters of each architectural class. Specifically, we varied the size of the embedding from 10 to 60 in intervals of 10. The number of units in the GRU/LSTM layers ranged from 8 to 384 in intervals of 8, and the number of units in the convolutional layer ranged from 4 to 192 in intervals of 4. For the convolutional layer, a size of 3 and a stride of 1 was used, which is based on the design principles from modern convolutional neural network [37]. No additional optimization was performed on the size or stride of the convolutional layers. In addition, no specific shape of the network topology was enforced.

#### 3.3 Bayesian Optimization of Hyperparameters


In order for Bayesian methods to be effective, a sufficient number of trials for different neural network design has to be performed. In practice, it has been recommended that a minimum of 10N trials



**Figure 1:** Illustration of the 4 architectural classes investigated, (a) GRU, (b) LSTM, (c) CNN-GRU and (d) CNN-LSTM. Number of units explored is indicated in parenthesis.

be performed, where N is the number of tunable hyperparameters. In our work, we performed 60 trials for each of the 4 architectural class. In addition, we manually seeded 6 initial designs for each class. Specifically, we used initial designs that had an embedding layer with 40, a convolution layer with 16 filters, and both LSTM/RNN layer with [8, 16, 32, 64, 128, 256] units.

In addition, because we are developing a general-purpose neural network design that can be re-used for a broad range of property prediction, it would not be feasible to include all conceivable training data to optimize the network design within the limits of available computing resources. Therefore, a subset of the datasets were used in the Bayesian optimization (see section 2.5 for details), and separate optimizations were performed for the Tox21 classification and the FreeSolv regression tasks.



**Figure 2:** Results of Bayesian optimization of the hyperparameters of the 4 architectural classes for the (a) Tox21 classification and (b) FreeSolv regression tasks.


The results of the Bayesian optimization across all 4 classes and 2 tasks are as indicated in Figure 2. For Tox21 classification, we observed that an additional convolutional layer between the embedding and RNN/LSTM layers improved model performance relative to their counterparts, and the best performing model was

em_size	#conv	#rnn1	#rnn2
50	192	224	384

**Table 2: Best CNN-GRU network design for the final SMILES2vec model.**

the CNN-LSTM class, with CNN-GRU trailing slightly behind. For FreeSolv regression, we observed that GRU-based networks outperform LSTM-based networks. Taking into considerations for generalization to other type of chemical properties, we selected the CNN-GRU architectural class for the remainder of this work. Then, we selected the best network design of this class, which is summarized in Table 2.

Lastly, because the Bayesian algorithm uses the validation metric as a means to optimize the network’s hyperparameters, there is a possibility that as one progresses, there may be overfitting towards the validation set. To determine the extent of overfitting, we examined the correlation between the validation metrics (whose validation set data would be changing during the Bayesian optimization) and the test metrics (whose test set data is fixed, and was never used in the Bayesian optimization). As illustrated in Figure 3, the correlation between validation and test metrics is 0.54 for the Tox21 dataset and 0.78 for the FreeSolv dataset. The lower correlation of the Tox21 dataset relative to the FreeSolv dataset may be explained by noting the AUC performance metric on which the optimization was performed, is not the same as the crossentropy loss function used for training the network.



**Figure 3: Correlation plot between validation and test metrics for all trials enumerated for the (a) Tox21 and (b) FreeSolv datasets.**

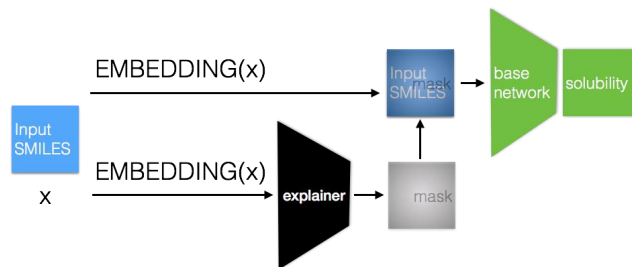
### 3.4 SMILES2vec Interpretation

To gain a better insight into the SMILES2vec model, we developed a method to gain some level of interpretability. Here, our objective is to identify the part(s) of the SMILES string that is responsible for the neural network’s decision.

Methods for explaining “black box” models exist [32], but most of these methods tend to require explicit combinatorial analysis. The approach we provide here provides insight into how the neural network analyzes the data, without combinatorially probing the input. This is achieved by training an explanation mask, whereby a separate explanation network learns to mask input data to produce near identical output as would be obtained from the original data.

### 3.5 Training the Explanation Mask

We train a neural network generated mask to identify the important characters of the input. The procedure is as follows: First, we use the final SMILES2vec model (Table 2) as the base network. Next, we construct another neural network to produce a mask over the input data, with the objective to train the mask such that the output of the base neural network remains the same but it masks as much data as possible. We freeze the base neural network, and we train the explanation mask end-to-end, as shown in Figure 4.



**Figure 4: Structure for training the explanation network. The SMILES input is passed through the embedding layer, then into the explainer. This produces a mask that is placed over the original embedding and sent through the pre-trained base model.**

With the weights of the base network frozen, the mask being learned will be specific to the SMILES2vec model. Each input will produce a different mask. To avoid the mask being trivial (completely uniform), we added two forms of regularization, a small L2 regularization, and we also penalized the mask for having high entropy (putting equal weight on all inputs). The overall loss function for a single element of each mini-batch is as follows:

$$Loss_i = ||f(SMILES_i, \theta) - Sol(SMILES_i)||_2 + 1e^{-6} ||MASK_i||_2 + 0.05H(MASK_i),$$

where  $f(SMILES_i, \theta)$  is the base neural network applied to the  $i$ th SMILES,  $Sol(SMILES_i)$  is the solubility,  $H$  is the entropy over the mask (normalized to sum to 1), and  $MASK_i$  is the vector of the calculated explanation mask at each entry in the input SMILES string.

The explanation network used to create the mask was a 20 layer residual network with SELU [26] activations. The padding was fixed such that the length of the input remained the same at each layer. The input to the network is the embedding of the SMILES string. The last layer is a 1D convolution of length 1, followed by a batch normalization, then by a softplus activation. We observed that the final batch normalization layer to be very important for trainability. This provides an output between 0 and infinity at each SMILES position. A mask output of 0 would prevent the base SMILES2vec from receiving that input character. A mask of  $\infty$  would cause the SMILES2vec to put more attention on that input character. We trained with Adam [25] until convergence. We started the learning rate at  $10^{-2}$  and divided by 10 as the training error plateaued, ultimately training down to  $10^{-6}$ .



SMILES	Solubility	Top-3 Chars
<b>c1ccncc1</b>	1.18	<b>c,c,n</b>
<b>O=Cc1ccco1</b>	-0.87	<b>C,c,o</b>
<b>Clc1ccc(cc1Cl)c2ccccc2</b>	-5.93	<b>c,c,c</b>
<b>Cc1c2ccccc2c(C)c3ccccc13</b>	-6.91	<b>c,c,c</b>

**Table 3: Sample SMILES entries, their predicted solubility value, and the top-3 most important characters are bolded.**

## 4 PERFORMANCE

In this section, we quantify the accuracy of interpretability on the solubility dataset. Then, we demonstrate generalizability of the SMILES2vec model by evaluating its performance on other datasets.


### 4.1 Interpreting Chemical Solubility

We demonstrate proof-of-concept for an interpretable SMILES2vec network using the ESOL solubility dataset. [42] Chemical solubility is a well-understood and simple chemical property where there is established first-principles knowledge. Briefly, parts of a chemical (i.e. functional groups) can typically be classified as either hydrophilic or hydrophobic. Hydrophilic “water-loving” groups, like alcohols, amines and carboxyl form strong interactions with water and increase the overall solubility of a compound, and they typically contain non-carbon atoms like nitrogen and oxygen. The reverse is true for hydrophobic “water-hating” groups, which tend to make chemicals more insoluble, and they typically are carbon-based chains/rings and halogens (chlorine, bromine, iodine).

We used the pre-trained the SMILES2vec base model, which attained a validation RMSE of 0.63. Solubility values are reported on the log10 scale, with less soluble compounds having more negative number. The mask outputs a normalized attention value that denotes the importance of a particular character in the network’s decision. For each SMILES string, we identified the top-3 characters (see Table 3 for examples). Then, we separated the dataset into soluble ( $\geq -1.0$ ) and insoluble ( $\leq -5.0$ ) compounds. Using established knowledge of chemical solubility to establish the ground truth, we expect that soluble compounds should have higher attention on the atoms O, N, and insoluble compounds to have higher attention on atoms C, F, Cl, Br, I. With this ground truth labeling in expected atoms, we computed the top-3 accuracy of SMILES2vec interpretability, which is 88%.

In addition, we also qualitatively examined the outputs of the masks by mapping the SMILES character to the corresponding atom(s) in the molecular structure, and examples are shown in Figure 5. For molecules with low solubility, the characters c, C, and Cl tend to receive more attention than others, which correspond to hydrophobic groups. In contrast, molecules with high solubility have attention focused on the characters O and N, which correspond to hydrophilic groups.

The localization of appropriate atoms on each functional group type depending on the chemical’s predicted solubility value indicates that SMILES2vec has learned representations that correspond to known chemistry concepts. Lastly, we emphasize that SMILES2vec has developed these representations without being provided any explicit chemical information. While chemical information is implicitly encoded in the SMILES string, no “decoding



**Figure 5: Colored circles of increasing darkness indicate the locations of increasing attention on the molecule. The explanation mask validates established knowledge by focusing on atoms of known hydrophobic and hydrophilic functional groups.**

solution” was provided to the network, neither was further feature engineering required. Our work therefore demonstrates the effectiveness of representation learning from raw data in the chemical sciences.

### 4.2 Generalization of SMILES2vec Models

Thus far, we have only evaluated the performance of SMILES2vec on the ESOL solubility dataset. Unlike solubility, the other 3 datasets (toxicity, activity, solvation energy) are more complex properties, for which no simple rule-based methods exist in the chemistry literature. Without the ability to generate ground-truth labels, quantifying the accuracy of the SMILES2vec interpretation is non-trivial, and is beyond the scope of this work. Nevertheless, the accuracy of the model’s predictions can still be evaluated.


We also note that the architectural optimization of SMILES2vec only included a small fraction of the 4 datasets identified for this work; the HIV dataset was not included, and 11 out of 12 toxicity tasks were not included. Hence, this section also determines how generalizable the Bayesian optimized network design will be to other chemical tasks.

First, we determined the effectiveness of generalizing SMILES2vec to the 3 remaining datasets. The following validation performance metrics were obtained: AUC of 0.80 for the full Tox21 dataset, AUC of 0.78 for the HIV dataset, and RMSE 1.4 kcal/mol for the FreeSolv dataset. Furthermore, we note that in all models, the difference between the validation and test metrics is small, further confirming the generalization of the model to compounds it has not seen either during the model training, or in the Bayesian hyperparameter optimization. Using a recently developed pre-training approach, [15] we were also able to improve the performance of SMILES2vec slightly, attaining AUC of 0.81 for the full Tox21 dataset, AUC of 0.80 for the HIV dataset, and RMSE 1.2 kcal/mol for the FreeSolv dataset.

Based on these results, we conclude that the Bayesian optimization of the network architectural design was effective in developing a general-purpose SMILES2vec network design for other chemical properties. We also note in recent literature there has been a trend towards using other “black box” approaches as a solution for network architecture design, for example using RNNs and reinforcement learning to optimize the design of a target neural network. [43] However, such methods typically require on the order of ~10,000 trials, which is much more than the ~500 trials used in our work.

In addition, given that the template of each architectural class was fixed, adaptive methods that automatically grow or shrink the neural network are also viable alternatives to network design. [35]

Next, we compare the performance of the best SMILES2vec model against contemporary deep neural networks that have reported results on the same datasets (Tox21, HIV, FreeSolv) that we have evaluated our model on. We compare against a typical MLP network that uses engineered features [42], a chemistry-specific molecular graph convolutional neural network [42], and Chemception, a deep CNN that uses images [17].



**Figure 6: Performance of SMILES2vec against contemporary deep neural networks trained on engineered features, image and graph data. For Tox21 and HIV, higher AUC is better. For FreeSolv, lower RMSE is better.**

The results are presented in Figure 6, and we use validation metrics to evaluate the quality of the model. In comparing the 4 methods, we observed that the standard MLP DL models that uses engineered features (fingerprints) performed the worst. SMILES2vec outperformed CNN models such as Chemception in classification tasks, but slightly underperformed in regression tasks. As indicated in our previous work, it is likely that the lack of atomic number information, which is not embedded in the SMILES format is responsible for its lower performance in predicting calculable physical properties. [17] In addition, SMILES2vec also outperformed first-principles models for computing solvation free energy (note: there are no first-principles models for computing toxicity/activity), which is especially noteworthy since SMILES2vec (and neural networks in general) can predict values much faster than traditional computational chemistry simulations, which typically require minutes to hours for each calculation. Against convolutional graphs, which is the current state-of-the-art for many chemical tasks, SMILES2vec either matches for classification tasks (Tox21: 0.81 vs 0.81, HIV: 0.80 vs 0.80) or outperforms for regression task (FreeSolv: 1.2 vs 1.3). Therefore, SMILES2vec is not only as accurate as the current state-of-the-art in chemistry DL models, but more importantly it is also an interpretable model.

## 5 CONCLUSION

In this paper, we develop SMILES2vec, the first general-purpose deep neural network that uses chemical text data (SMILES) for predicting chemical property, with an explanation mask that improves interpretability. By performing extensive Bayesian optimization experiments, we identified a specific CNN-GRU neural

network architecture is effective in predicting a wide range of properties. SMILES2vec achieved a validation AUC of 0.81 and 0.80 for Tox21 toxicity and HIV activity prediction respectively, and a validation RMSE of 1.2 kcal/mol and 0.63 for solvation energy and solubility. Using the solubility dataset as an illustration of SMILES2vec interpretability, we construct explanation masks that indicate SMILES2vec localizes on specific characters in hydrophilic or hydrophobic groups, with a top-3 accuracy of 88%. Identification of such functional groups and their relationship to chemical solubility is a key first-principles concept in chemistry, which SMILES2vec was able to discover on its own. Compared to other DL models, SMILES2vec’s accuracy outperforms the typical MLP DL models that uses engineered features as input. Against the current state-of-the-art (convolutional graph networks), SMILES2vec outperforms on regression tasks and matches on classification tasks. These results indicate that SMILES2vec can accurately predict a broad range of properties and learn technically accurate chemical concepts, which suggest that it can be used as an interpretable tool for the future of deep learning driven chemical design.

## ACKNOWLEDGMENTS

The authors would like to thank Dr. Nathan Baker for helpful discussions. This work is supported by the following PNNL LDRD programs: Pauling Postdoctoral Fellowship and Deep Learning for Scientific Discovery Agile Investment.

## REFERENCES

- [1] Martın Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. TensorFlow: A System for Large-Scale Machine Learning.. In *OSDI*, Vol. 16. 265–283.
- [2] Esben Jannik Bjerrum. 2017. SMILES Enumeration as Data Augmentation for Neural Network Modeling of Molecules. *arXiv preprint arXiv:1703.07076* (2017).
- [3] Peter Buchwald and Nicholas Bodor. 2002. Computer-aided drug design: the role of quantitative structure–property, structure–activity and structure–metabolism relationships (QSPR, QSAR, QSMR). *Drugs Future* 27, 6 (2002), 577–588.
- [4] Artem Cherkasov, Eugene N Muratov, Denis Fourches, Alexandre Varnek, Igor I Baskin, Mark Cronin, John Dearden, Paola Gramatica, Yvonne C Martin, Roberto Todeschini, et al. 2014. QSAR modeling: where have you been? Where are you going to? *Journal of medicinal chemistry* 57, 12 (2014), 4977–5010.
- [5] Sharan Chetlur, Cliff Woolley, Philippe Vandermersch, Jonathan Cohen, John Tran, Bryan Catanzaro, and Evan Shelhamer. 2014. cudnn: Efficient primitives for deep learning. *arXiv preprint arXiv:1410.0759* (2014).
- [6] Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259* (2014).
- [7] John D Chodera, David L Mobley, Michael R Shirts, Richard W Dixon, Kim Branson, and Vijay S Pande. 2011. Alchemical free energy methods for drug discovery: progress and challenges. *Current opinion in structural biology* 21, 2 (2011), 150–160.
- [8] François Chollet et al. 2015. Keras. (2015).
- [9] George E Dahl, Navdeep Jaitly, and Ruslan Salakhutdinov. 2014. Multi-task neural networks for QSAR predictions. *arXiv preprint arXiv:1406.1231* (2014).
- [10] Ian Dewancker, Michael McCourt, Scott Clark, Patrick Hayes, Alexandra Johnson, and George Ke. 2016. Evaluation System for a Bayesian Optimization Service. *arXiv preprint arXiv:1605.06170* (2016).
- [11] Li Di, Paul V Fish, and Takashi Mano. 2012. Bridging solubility between drug discovery and development. *Drug Discovery Today* 17, 9–10 (2012), 486–495.
- [12] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. 2015. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in neural information processing systems*. 2224–2232.
- [13] Anna Gaulton, Louisa J Bellis, A Patricia Bento, Jon Chambers, Mark Davies, Anne Hersey, Yvonne Light, Shaun McGlinchey, David Michalovich, Bissan Al-Lazikani, et al. 2011. ChEMBL: a large-scale bioactivity database for drug discovery. *Nucleic acids research* 40, D1 (2011), D1100–D1107.

- [14] Garrett B Goh, Nathan O Hodas, and Abhinav Vishnu. 2017. Deep learning for computational chemistry. *Journal of Computational Chemistry* (2017).
- [15] Garrett B Goh, Charles Siegel, Abhinav Vishnu, and Nathan O Hodas. 2017. ChemNet: A Transferable and Generalizable Deep Neural Network for Small-Molecule Property Prediction. *arXiv preprint arXiv:1712.02734* (2017).
- [16] Garrett B Goh, Charles Siegel, Abhinav Vishnu, Nathan O Hodas, and Nathan Baker. 2017. Chemception: A Deep Neural Network with Minimal Chemistry Knowledge Matches the Performance of Expert-developed QSAR/QSPR Models. *arXiv preprint arXiv:1706.06689* (2017).
- [17] Garrett B Goh, Charles Siegel, Abhinav Vishnu, Nathan O Hodas, and Nathan Baker. 2017. How Much Chemistry Does a Deep Neural Network Need to Know to Make Accurate Predictions? *arXiv preprint arXiv:1710.02238* (2017).
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*. 1026–1034.
- [19] G Hinton, N Srivastava, and K Swersky. 2012. RMSProp: Divide the gradient by a running average of its recent magnitude. *Neural networks for machine learning, Coursera lecture 6e* (2012).
- [20] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [21] Tyler B Hughes, Na Le Dang, Grover P Miller, and S Joshua Swamidass. 2016. Modeling Reactivity to Biological Macromolecules with a Deep Multitask Network. *ACS central science* 2, 8 (2016), 529–537.
- [22] Stanisaw Jastrzebski, Damian Lesniak, and Wojciech Marian Czarnecki. 2016. Learning to SMILE (S). *arXiv preprint arXiv:1602.06289* (2016).
- [23] William L Jorgensen. 2004. The many roles of computation in drug discovery. *Science* 303, 5665 (2004), 1813–1818.
- [24] Steven Kearnes, Kevin McCloskey, Marc Berndl, Vijay Pande, and Patrick Riley. 2016. Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design* 30, 8 (2016), 595–608.
- [25] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [26] Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. 2017. Self-normalizing neural networks. In *Advances in Neural Information Processing Systems*. 972–981.
- [27] Naomi L Kruhlik, Joseph F Contrera, R Daniel Benz, and Edwin J Matthews. 2007. Progress in QSAR toxicity screening of pharmaceutical impurities and other FDA regulated products. *Advanced drug delivery reviews* 59, 1 (2007), 43–55.
- [28] Andreas Mayr, Günter Klambauer, Thomas Unterthiner, and Sepp Hochreiter. 2016. DeepTox: toxicity prediction using deep learning. *Frontiers in Environmental Science* 3 (2016), 80.
- [29] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [30] John R Platt. 1947. Influence of neighbor bonds on additive bond properties in paraffins. *The Journal of Chemical Physics* 15, 6 (1947), 419–420.
- [31] Bharath Ramsundar, Steven Kearnes, Patrick Riley, Dale Webster, David Konerding, and Vijay Pande. 2015. Massively multitask networks for drug discovery. *arXiv preprint arXiv:1502.02072* (2015).
- [32] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. Why should I trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1135–1144.
- [33] David Rogers and Mathew Hahn. 2010. Extended-connectivity fingerprints. *Journal of chemical information and modeling* 50, 5 (2010), 742–754.
- [34] David E Shaw, JP Grossman, Joseph A Bank, Brannon Batson, J Adam Butts, Jack C Chao, Martin M Deneroff, Ron O Dror, Amos Even, Christopher H Fenton, et al. 2014. Anton 2: raising the bar for performance and programmability in a special-purpose molecular dynamics supercomputer. In *Proceedings of the international conference for high performance computing, networking, storage and analysis*. IEEE Press, 41–53.
- [35] Charles Siegel, Jeff Daily, and Abhinav Vishnu. 2016. Adaptive neuron apoptosis for accelerating deep learning on large scale systems. In *Big Data (Big Data), 2016 IEEE International Conference on*. IEEE, 753–762.
- [36] John E Stone, David J Hardy, Ivan S Ufimtsev, and Klaus Schulten. 2010. GPU-accelerated molecular modeling coming of age. *Journal of Molecular Graphics and Modelling* 29, 2 (2010), 116–125.
- [37] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1–9.
- [38] Roberto Todeschini and Viviana Consonni. 2008. *Handbook of molecular descriptors*. Vol. 11. John Wiley & Sons.
- [39] Izhar Wallach, Michael Dzamba, and Abraham Heifets. 2015. AtomNet: a deep convolutional neural network for bioactivity prediction in structure-based drug discovery. *arXiv preprint arXiv:1510.02855* (2015).
- [40] David Weininger. 1988. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *Journal of chemical information and computer sciences* 28, 1 (1988), 31–36.
- [41] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144* (2016).
- [42] Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. 2017. MoleculeNet: A Benchmark for Molecular Machine Learning. *arXiv preprint arXiv:1703.00564* (2017).
- [43] Barret Zoph and Quoc V Le. 2016. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578* (2016).