
BERT Learns (and Teaches) Chemistry

Josh Payne^{1,2}, Mario Srouji^{1,2}, and Dian Ang Yap^{1,2}

¹Department of Computer Science, Stanford University
²{joshp007, msrouji, dianang7}@stanford.edu

Abstract

Modern computational organic chemistry is becoming increasingly data-driven. There remain a large number of important unsolved problems in this area such as product prediction given reactants, drug discovery, and metric-optimized molecule synthesis, but efforts to solve these problems using machine learning have also increased in recent years. In this work, we propose the use of attention to study functional groups and other property-impacting molecular substructures from a data-driven perspective, using an transformer-based model (BERT) on datasets of string representations of molecules and analyzing the behavior of its attention heads. We then apply the representations of functional groups and atoms learned by the model to tackle problems of toxicity, solubility, drug-likeness, and synthesis accessibility on smaller datasets using the learned representations as features for graph convolution and attention models on the graph structure of molecules, as well as fine-tuning of BERT. Finally, we propose the use of attention visualization as a helpful tool for chemistry practitioners and students to quickly identify important substructures in various chemical properties.

1 Introduction

There are several problems in organic chemistry, such as the prediction of chemical properties of unknown molecules, that are very difficult and remain largely unsolved. Handcrafting approaches are not scalable as the number of reported organic reactions constantly increases and significant time investment is needed to keep up with the literature. However, certain heuristics to discover new molecules have proven useful for these tasks, such as the notion of a functional group. Functional groups are sets of connected atoms that determine properties and reactivity of a parent molecule. They are a cornerstone of modern organic chemistry and chemical nomenclature. There are many scientific papers that focus on properties and reactivity of various functional groups, however there is little attention paid to the study of functional groups from a data-driven point of view.

In modern applications such as drug discovery, the search space is estimated to be in the range of 10^{23} to 10^{60} molecules (20); however, functional groups and property-impacting substructures can greatly assist in restricting the size of this space, as they indicate the presence of specific chemical properties in molecules. Understanding the influence of these functional groups in relation with each other is also a crucial step in understanding molecular properties such as stability and reactivity, besides providing insights into activating/deactivating properties when we seek to synthesize new molecules. They remain a widely researched topic in chemistry, ranging from understanding activating/deactivating groups, steric hindrance, and resonance in large molecules.

Currently, the majority of theoretical studies utilize functional groups as a rule-based definition to categorize compounds into classes, allowing chemists to perform large-scale automated chemical classification based on a structure-based chemical taxonomy. Additionally, while machine learning has been applied to predict reactions, and to generate new molecular graphs, little work has been done to either leverage functional groups to perform predictions or to interpret the deep learning models to identify functional groups (hence allowing the groups to emerge from the data as opposed to from heuristics).

In this work we propose a method to perform automated identification of important molecular substructures without any chemistry domain expertise, hard-coding of structures, heuristic methods, pre-selection approaches, or restricted definitions. We instead detect important functional groups of atoms in molecules by allowing them to emerge from the data. In this way, we aim to learn these rules of chemistry directly. We do this through the use of attention mechanisms on the SMILES strings of molecules, by setting up the task of identifying important substructures in BERT’s pretraining procedure and a following task of fine-tuning BERT for regression on solubility, drug-likeness, and synthesis accessibility, in order to identify important substructures of molecules.

2 Related Work

Language Modeling

BERT (Bidirectional Encoder Representations from Transformers) (5) was demonstrated to outperform previous language models on a variety of tasks with high efficiency. Other works also explore multi-task learning on language models that leverage large amounts of cross-task data, which also exhibit from a regularization effect that leads to more general representations to help adapt to new tasks and domains (18). Moreover, recent work and tools have shown that interpretability is tractable with BERT (4; 12).

BERT builds upon the encoder structure of the encoder-decoder architecture of transformer, which uses multi-head self attention (2; 3). Different techniques have been developed to better interpret attention maps, such as attention matrix heatmaps (2; 27; 25) and bipartite graph representations (17; 30).

Multitask and Meta-Learning

Although the field of meta learning is relatively nascent, there are still several technical approaches that have been utilized on similar problems that we believe can be applied to computational chemistry. For instance, Ramsundar et. al. (23) were able to demonstrate how increasing the amount of data and tasks can greatly increase performance when discovering novel drugs; this parallels our initial hypothesis with discovering functional groups. Further, Ren et. al. (24) augmented prototypical networks for semi-supervised classification; we believe such an approach may also help us with classifying and grouping functional groups and molecules. Finally, Garcia et. al. (8) applied few shot learning to GNNs; as molecules and reactions can be modeled using graphs and hypergraphs, we believe a similar approach may be useful here.

Using Machine Learning for Chemical Inference Tasks

Previous work leveraged a transformer model to perform chemical reaction prediction (28). Given reactants and reagents, they predict the products, and similar to other work, they treat reaction prediction as a machine translation problem between the SMILES strings. They show that a multi-head attention Molecular Transformer model outperforms all algorithms in the literature, requiring no handcrafted rules. They did not however leverage functional groups in their predictions, or attempt to interpret the model to detect groupings of atoms, and their effects.

Another work by You et al. (35) proposes the Graph Convolutional Policy Network (GCPN), an approach to generate molecules where the generation process can be guided towards specified desired objectives, while restricting the output space based on underlying chemical rules. Graph representation learning is used to obtain vector representations of the state of generated graphs, adversarial loss is used as reward to incorporate prior knowledge specified by a dataset of example molecules, and the entire model is trained end-to-end in a reinforcement learning framework.

One slightly related work creates molecule Deep Q-Networks (MolDQN), for molecule optimization by combining domain knowledge of chemistry, and reinforcement learning techniques (36). They define modifications on molecules to ensure chemical validity. MolDQN achieves comparable or better performance against several other recently published algorithms for benchmark molecular optimization tasks. However, they argue that many of these tasks are not representative of real optimization problems in drug discovery. They extend their model with multi-objective reinforcement learning, to maximize drug-likeness while maintaining similarity to the original molecule. In both of the above works, we hypothesize that the attention head behavior can be used to make the reward function more dense, decreasing training time.

Regression tasks on chemical attributes such as toxicity, solubility (logP), drug-likeness (QED), and synthesis accessibility (SAS) have been proposed by several works. Graph Neural Networks (namely, Graph Convolutional Networks and Graph Attention Networks) have been shown to be effective on toxicity tasks such as the Tox21 dataset (15; 33), and recently, pretraining procedures for graph neural networks have been shown to be effective on chemical inference (13). Explainability and identification of important substructures on molecule graphs have also been recently proposed in (34) using GATs. However, these graph-based machine learning approaches are inherently limited in their depth, as too many layers will cause oversmoothing. This issue doesn't exist with deep language models such as BERT, which is built using multiple transformers and hundreds of millions of parameters.

However, there exists the issue of learning the encoding rules built into SMILES strings, which are formally defined in Section 3. One problem that arises here is that two unique SMILES strings representing the same molecule can have different embeddings. How can a language model be trained to recognize isomorphism between unique, yet equivalently mapping, SMILES strings, essentially learning a truer representation of the molecule the SMILES string encodes? One work proposes a very elaborate and hand-designed architecture that performs well on regression tasks using SMILES string (1). However, this architecture, while achieving state-of-the-art results, is not a clean solution that allows the encoding rules to arise from the data or training procedure.

Detecting Functional Groups within Molecules

There currently exists a small amount of work on identifying functional groups via an automated process, as in (6). The approach of this paper is based on processing hetero-atoms and their environment with the addition of some other functionalities, like multiple carbon-carbon bonds. However while this is a step forward from manual curation, the algorithm involves an iterative process of marking and merging of atoms based on pre-defined heuristics. This inherently restricts the ability to learn functional groups through data, or through chemical reaction properties. Additionally, the work admits that this algorithm does not provide an ultimate definition of functional groups, as every chemist has a slightly different understanding about what a functional group is. For example, in their present algorithm they restricted their definition only to classical acetal, thioacetal or aminal centers, and did not consider other similar systems, like alpha-substituted carbonyls, or similar bonds. Hence the method they used also suffers from a strict formulation of what constitutes a functional group within a molecule.

Another state-of-the-art program for identification of functional groups is called checkmol (9). Checkmol is a command-line utility program which reads molecular structure files in different formats, and analyzes the input molecule for the presence of various functional groups and structural elements, from a predefined set of 204 functionalities (through pattern matching). The philosophy of checkmol is slightly different from that of their approach, as checkmol defines functional groups in a hierarchical manner. In summary, both approaches had very similar results for the functional groups detected. This shows that systems based on a well selected list of substructures (as apparently checkmol is) can provide useful information about the functional group composition of general molecular data-sets, however both approaches involved expert knowledge in the field of chemistry in order to either manually curate a selection of functional groups, or to create a strictly defined heuristic to detect functional groups iteratively.

There have been a few machine learning-based approaches in identifying functional groups from data. Last month, Pope et al. (21) released a paper outlining a GCN-based approach to detecting functional groups in data. This work was an extension of their 2019 paper on explainability methods for GCNs (22) and proposed the application of contrastive gradient-based saliency maps, class activation mapping (CAM), and excitation backpropagation, previously proposed for convolutional neural networks, to the graph convolutional neural network domain. However, GCNs are foundationally less expressive than many modern sequence models, as their depth is restricted by the number of layers they can have before oversmoothing occurs. Additionally, they are not as readily explainable as sequence models with attention layers.

3 Approach


3.1 Datasets

Some molecules are more readily synthesizable than others; that is, the chemical properties present in synthesizable molecules may provide a stronger indication of functional groups. The first dataset we analyze is a set of 250,000 readily synthesizable molecules collected from the ZINC15 database¹. This dataset contains the SMILES string for each of these molecules, as well as the logP, QED, and SAS values. The term SMILES (Simplified Molecular Input Line Entry System) refers to a string notation for encoding molecular structures (31). In order to more easily perform computation and analysis on molecules, David Weininger developed the SMILES (Simplified Molecular Input Line Entry System) specification at the USEPA Mid-Continent Ecology Division Laboratory in Duluth in the 1980s.

Typically, numerous equally valid SMILES strings can be written for a molecule. SMILES contains the same information as might be found in an extended connection table. The primary reason SMILES is more useful than a connection table is that it is a linguistic construct, rather than a computer data structure. SMILES is a true language, with a simple vocabulary (atom and bond symbols) and only a few grammar rules. SMILES notation consists of a series of characters containing no spaces. Hydrogen atoms may be omitted (hydrogen-suppressed graphs) or included (hydrogen-complete graphs).

In terms of a graph-based computational procedure, SMILES is a string obtained by printing the symbol nodes encountered in a depth-first tree traversal of a chemical graph. The chemical graph is first trimmed to remove hydrogen atoms and cycles are broken to turn it into a spanning tree. Where cycles have been broken, numeric suffix labels are included to indicate the connected nodes. Parentheses are used to indicate points of branching on the tree. Below is an example SMILES string and corresponding molecule:

CC(C)(C)c1ccc2occc(CC(=O)Nc3ccccc3F)c2c1



After studying chemical properties with attention in the ZINC15 250k dataset, we want to apply our new representations to a smaller, more difficult task. The Tox21 dataset is a multitask dataset of 12 unique toxicity tasks, each a binary categorical task. It contains roughly 7,000 molecules and typical featurizations include atomic number, aromaticity, donor status, and acceptor status.

3.2 Neural Machine Translation and Language Modeling

Neural machine translation was proposed by Kalchbrenner and Blunsom (2013), Sutskever et al. (2014) and Cho et al. (2014b). An encoder neural network reads and encodes a source sentence into a fixed-length vector. A decoder then outputs a translation from the encoded vector. The whole encoder-decoder system is jointly trained to maximize the probability of a correct translation given a source sentence.

Recurrent neural networks (RNN) (19), and long short-term memory (LSTM) (11) have been established as the state of the art approaches in sequence modeling, and are commonly used in Neural Machine translation. Recurrent models factor computation along the symbol positions of the input and output sequences, by aligning the positions to steps in computation time. They generate a sequence of hidden states h_t , as a function of the previous hidden state h_{t-1} , and the input for position t . The issue is that a neural network needs to be able to compress all the necessary information of a source sentence into a fixed-length vector. This may make it difficult for the neural network to cope with

¹Link to the ZINC15 250k dataset: https://github.com/aspuru-guzik-group/chemical_vae/tree/master/models/zinc

long sentences, especially those that are longer than the sentences in the training corpus. Cho et al. (2014b) showed that the performance of a basic encoder–decoder deteriorates rapidly as the length of an input sentence increases.

In order to address this issue, attention mechanisms (2) were introduced as an extension to the encoder–decoder model. The generic attention takes in three values: a value (V), a key (K), and a query (Q). From these values, it applies a standard dot-product attention, as shown in Equation 1, with a scaling factor of $\frac{1}{\sqrt{d_k}}$ added to ensure that the value of the dot product doesn’t grow too large in magnitude with respect to d_k , the dimension of the key.

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V \quad (1)$$

Attention models also have high interpretability, allowing the researcher to observe what the attention heads are "attending" to by investigating the values of the attention heads as inputs are fed through, as done in BERT.

3.3 Multi-Task Learning

In this work, we propose a tripartite approach to studying and understanding functional groups and applying these findings to studying chemical reactions and properties. These datasets will be investigated after we’ve learned to learn representations of functional groups from the larger datasets.

To create meaningful representations of atomic tokens in the context of molecules and reaction networks, we intend to train BERT on several large reaction datasets, including the ZINC15 250k dataset. A multi-task learning approach will be used: by learning “masked” molecules in string representations, relations learned between molecular string representations will hopefully pave the way for downstream tasks such as link prediction, functional group identification and generative tasks. Moreover, we hope to achieve SoA accuracy on reaction prediction tasks with these datasets.

We want to explore the explainability and interpretability of this model and use it to identify functional groups. Our goal is to be able to use these predictions to validate the model’s interpretability using existing knowledge about functional groups as well as identify new, useful functional groups for chemists. Our hypothesis is that an analysis of the behaviour of BERT’s attention heads will give indications of certain atoms which form functional groups. Interpretability of multi-task learning and meta-learning remains an intriguing field that would pave better understanding with more research.

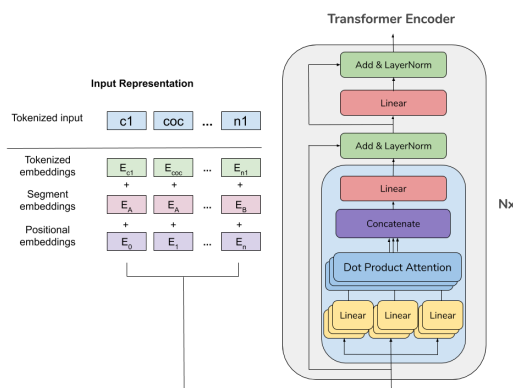


Figure 1: Visualization of BERT trained on SMILES strings. SMILES sequences are tokenized into functional groups, and the encoded embeddings with positional and token embeddings are fed in as input to BERT, the encoder structure of transformer. The light blue box denote attention mechanisms used in BERT.

Finally, we would like to transfer this knowledge to tasks on smaller, more difficult datasets, such as the Tox21 (14) and MUV (26) multitask datasets. We intend to compare results of using learned functional group knowledge explicitly, features computed by BERT, combined with various language and graph learning models (LSTM (11), BERT (5), Transformer (32), GCN (15), GraphSAGE (10),

and Graph Attention Networks (33), as molecules may be represented as SMILES strings and graphs equivalently). We’d also like to explore meta-learning techniques such as MAML (7) for improved performance in low-data settings.

4 Experimental Setup

4.1 Pre-analysis using n -grams

Before training our BERT model, we performed a pre-analysis of the ZINC15 250k dataset using n -gram (or length n subsequence) frequency evaluation in an attempt to find the most common atomic motifs within the molecules. Our hypothesis was that common motifs in synthesizable molecules correspond often to functional groups, as their synthesizability lends them nice reactive properties. An n -gram, in this context, is defined as a continuous sequence of atoms and bonds in the SMILES string representation of the molecule. For each molecule, we gathered all possible contiguous sub-sequences (of all possible lengths based on the SMILES string), by varying our n -gram "window" size. We then created a histogram of how frequently each SMILES sub-sequence (gathered through n -grams) appears across all of the molecules in the dataset. What we found is that common functional groups indeed appear with high frequency in this dataset; for instance, the 6-atom benzene ring appears in the top 50 most frequent n -grams.

4.2 BERT learns Chemistry

From the raw datasets, we tokenize using the regular expression proposed by (29). We also add tokens such as [UNK], [SEP], [MASK], [PAD] to denote special tokens such as unknown tokens, separation, mask and padding tokens. Similar to BERT trained on natural language models that mask out 15% of the tokens, we mask out 15% of the tokens in each SMILES string with a maximum sequence length of 128, which the model then learns to predict the masked out atoms and functional groups. This is equivalent to the pretraining step where by learning to predict the masked token, the molecular BERT could learn transferable tasks such as property prediction.

We employ a bidirectional context to learn contextual representations on molecular corpus. We also encode positional embeddings of length 512 on the tokenized inputs to learn relative positioning, and employ layer normalization and residuals to propagate signals through the deep 12-layer network. feed-forward sublayers within each layer computes non-linear hierarchical features, with 10% dropout probability to improve generalization. We use 12 attention heads for the self-attention layers, which are then extracted for interpretable attention visualization.

Training. We train the model with a batch size of 128 on a learning rate of 2×10^{-5} for 500,000 steps, with a linear warmup followed by Adam optimizer, where the initial linear warmup addresses the large variance in early learning stages (16). This was conducted on a Google TPU over the course of 24 hours. At this point, the loss had converged to nearly zero.

4.3 Analysis of BERT’s Attention

Leveraging the self-attention mechanism, we visualize attention structures of BERT trained on chemical molecular representations such as SMILES, which could uncover new relationships on how different functional groups within a molecule attend and influence each other.

The BERT attention layer takes in an input of vector sequence $h = [h_1, h_2, \dots, h_n]$ corresponding to the n tokenized functional groups of the SMILES sequence. Each vector h_i is transformed into query, key and value vectors (q_i, k_i, v_i) through separate affine transformations. The head computes attention weights between all pairs of tokens as softmax-normalized dot products between the query and key vectors. For token i and for value vector j ,

$$\alpha_{ij} = \frac{\exp(q_i^T k_j)}{\sum_{n=1}^n \exp(q_i^T k_l)} \quad (2)$$

The output o_i of the attention head i is a weighted sum of the value vectors, which represents how important every other tokenized functional group is when producing the next representation for the



(a) LogP, or solubility, is highly impacted by the length of carbon chains. We noted that several attention heads of BERT fine-tuned on logP value in the 10th layer attend heavily to carbon chains in this molecule.

(b) The attention heads of BERT pretraining attention does not pay special attention to carbon chains or other particular substructure.

Figure 2: The attention heads focus more heavily on the carbon chains; a carbon chain’s length can help determine solubility.

current tokenized functional group.

$$o_i = \sum_{j=1}^n \alpha_{ij} v_j \quad (3)$$

While previous work on visualizing attention were mostly on encoder-decoder models, we tweaked the classic transformer visualization to visualize attention maps on the encoder-only BERT. For a given transformer layer, we could visualize how a molecule, in its SMILES representation, attends to itself or to another molecule. Moreover, since parameters are not shared across the attention heads, each attention head is capable of producing a unique attention pattern. The attention heads are capable of capturing relative positional patterns and specific structural (lexical) patterns across atoms in molecules.

5 Attention-based Analysis of Chemically-Attributed Molecular Substructures

For a molecule m , we define atom a_m as atom a in molecule m . We would like to investigate which atom a_m in molecule m would react with other molecules m' , where $m' \in \{M \setminus m\}$, i.e. to identify the most likely site of m where other molecules will react with. This part is done after pretraining the BERT structure on the chemical molecule datasets in an unsupervised manner by filling in the masked tokens as illustrated in the pretraining step above. Define the function $f_\theta : (M_1, M_2) \rightarrow \mathbb{R}^n$ denote the attention from molecule $m \in M_1$ to $m' \in M_2$ parameterized by θ , which are BERT’s parameters. We have the extracted attention a_m as

$$\tilde{a}_m = \sum_{m' \in \{M \setminus m\}} \sum_{a_i \in m'} f_\theta(a_i, a_m) \quad (4)$$

$$a_m = \frac{\tilde{a}_m - \min(a_m)}{\max(a_m) - \min(a_m)} \quad (5)$$

Surprisingly, despite only being trained on SMILES string representation and learning the structure of sequence of tokenized atoms, without any data on reactions, BERT’s attention heads learns to put more weight to sites which are, in reality, active functional groups that act as a site for reaction. Consider the example in Figure 3:

Despite only learning the structural representation of the molecule, it learns the active reaction sites as potential functional groups. In the example above, we note that the carbon in aspartate (C) attends less to the α -ketoglutarate molecule as a whole, whereas the reaction site of $[\text{NH}_3^+]$ attends more strongly to the whole molecule of α -ketoglutarate. Other visualizations are illustrated as per Figure 4.

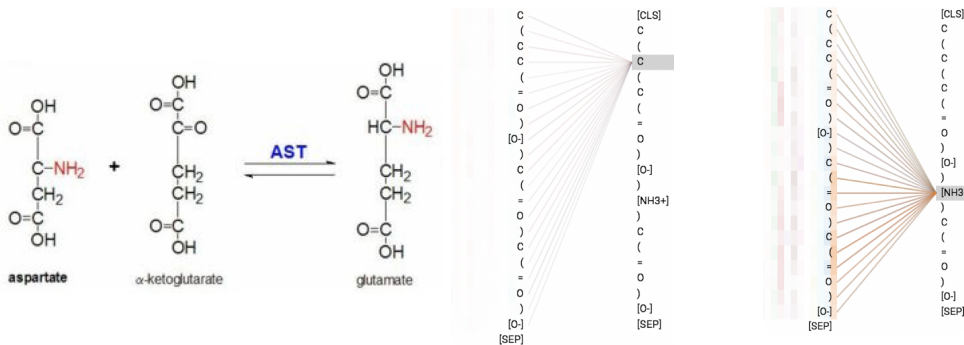


Figure 3: Left: illustration of chemical reaction on how aspartate reacts with α -ketoglutarate to form glutamate. Right: Visualization of layer 3 of attention heads in BERT on different molecules in aspartate.




Figure 4: Molecular Group discovery of different chemical compounds.

6 Multi-task learning for downstream tasks

6.1 Generalizing to the graph domain

After pretraining BERT in an unsupervised manner on ZINC250k, we want to see whether the chemical intuition BERT has picked up is ready to be transferred to other datasets and tasks. In this investigation, we considered the Tox21 dataset, as well as two standard models for inference on it, Graph Convolutional Networks and Graph Attention Networks, comparing different types of featurizations.

6.1.1 Graph Convolutional Networks

Letting R be the matrix of features for vertices r_v , A be the adjacency matrix of G , D be the diagonal matrix where $D_{i,i} = \sum_j (A + I)_{i,j}$, and $f(R, A)$ be the function we would like to find that we would like to learn, the $l + 1^{\text{st}}$ graph convolutional layer is defined as

$$H^{(l+1)} = \sigma\left(D^{-\frac{1}{2}}(A + I)D^{-\frac{1}{2}}H^{(l)}W^{(l)}\right)$$

where $H^{(0)} = R$. This model captures the $(l + 1)$ -hop neighborhood topology of a vertex.

6.1.2 Graph Attention Networks

We will describe a single graph attention layer, as it is the sole layer utilized throughout.

The input to our layer is a set of node features, $h = [h_1, h_2, \dots, h_N], h_i \in R^F$, where N is the number of nodes, and F is the number of features in each node. The layer produces a new set of node features after the transformation.

In order to obtain sufficient expressive power to transform the input features into higher-level features, at least one learnable linear transformation is required. Then self-attention is performed on the nodes to compute attention coefficients $e_{ij} = a(Wh_i, Wh_j)$ that indicate the importance of a given node j 's features to node i . In its most general formulation, the model allows every node to attend on every

other node, dropping all structural information. The graph structure is injected into the mechanism by performing masked attention—only computing coefficients for nodes based on their neighborhood in the graph. To make coefficients easily comparable across different nodes, they are normalized across all choices of node using the softmax function. Once obtained, the normalized attention coefficients are used to compute a linear combination of the features corresponding to them, to serve as the final output features for every node.

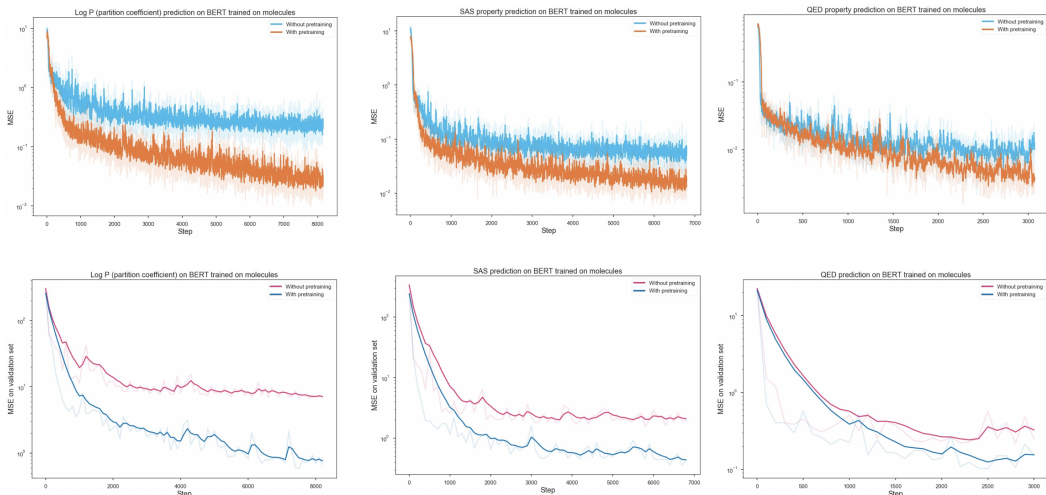


Figure 5: Illustration of BERT used for multi-task learning after pretraining. We test the Zinc250k dataset after pretraining on 3 different task: LogP, the partition coefficient (left), SAS, synthetic accessibility score (middle), QED, the quantitative estimation of drug-likeness (right). Top: training set MSE, bottom: test set MSE.

Table 1: Comparison of effectiveness of baseline, baseline concatenated with BERT embeddings, and with BERT embeddings only evaluated on GCN and GAT models.

Model	Validation Accuracy (%)	
	GCN	GAT
Baseline (canonical atom featurization, CAF)	0.815	0.826
CAF concatenated with BERT embeddings	0.817	0.827
BERT embeddings only	0.794	0.795
One-hot encoding of atomic number	0.790	0.788
Random featurization	0.678	0.673

From these results, we found that the embeddings that BERT learns for molecules are not useful in improving upon the state of the art; what it learns seems to be a subset of the information captured by the canonical atom featurizer. One potential detractor from BERT’s performance is that there exist molecules in Tox21 with atoms that do not appear in BERT’s vocabulary. These are not very many relatively though and should not have a huge impact.

6.2 Fine-tuning BERT for regression tasks

Once we pretrained BERT, we decided to fine-tune it with respect to the three values in the ZINC15 dataset. BERT was able to perform this task fairly well, and we found that, intuitively, BERT learning the task with random initial weights did not perform as well as BERT learning the task with pretraining (in both training and validation for all three tasks, as shown in Figure 5).

We also felt that it would be interesting to visualize BERT’s attention heads on molecules after it had been fine-tuned to perform a specific chemical inference task. However, we found that the resulting attention differed very little on the whole from task to task, even from pretraining as well. The attention heads displayed seemingly identical behavior for each of the layers across each of the tasks,

with the only standout exception being the attention heads for the solubility task in layers 10 and 11. These are the attention heads showcased in Figure 2. Some possible causes of this may include the low value range of the QED and SAS tasks (if we unnormalized the regression values, would the gradients be larger, hence causing more obvious change in behavior in the attention heads?) or that BERT’s fine-tuning step learns a local optima that is close to the pretraining weights but does not answer anything useful about chemistry. However, the attention heads on the solubility task seemed to behave intuitively, showing that there is a glimmer of hope in this line of investigation. We hope that this is the case, given that a working model like this would be an incredibly helpful tool for chemists.

7 Future Work and Extensions

We plan to continue this work using a larger dataset of 96,000,000 molecules gathered from the PubMed dataset², as we’ve observed that language models like BERT seem to thrive in settings with an abundance of data. Two challenges here are the sheer computational resources required to wield and train on such a massive dataset, and the determination of useful molecules that do not add to the noise in the signal.

Additionally, we would like to design a loss function for BERT that penalizes distances between embeddings of SMILES strings that represent the same molecule yet are unique. We hope that this will force BERT to learn the inherent graphical meaning of the SMILES strings and hence to start to pick up all of the chemical inference that can be gathered based on connectivity. For this, we will need to augment a smaller dataset, generating multiple unique SMILES strings for the same molecule for each of the molecules in the dataset. At this moment, we are training BERT on such a dataset for future investigation.

Finally, we would like to identify potential flaws in our approach to analyzing chemical property-bearing molecular substructures using fine-tuning. We feel that if implemented correctly, this could be a very useful tool for chemists and could greatly assist drug discovery, reaction prediction, and other tasks.

8 Discussion and Conclusion

In this work, we proposed the analysis of BERT and its attention to study functional groups and other property-impacting molecular substructures from a data-driven perspective on datasets of SMILES representations of molecules. We then applied the representations of functional groups and atoms learned by the model to tackle problems of toxicity, solubility, drug-likeness, and synthesis accessibility on smaller datasets using the learned representations as features for graph convolution and attention models on the graph structure of molecules, as well as fine-tuning of BERT. Finally, we proposed the use of attention visualization as a helpful tool for chemistry practitioners and students to quickly identify important substructures in various chemical properties. While we haven’t defined the state of the art with any of our results or designed any breakthrough tools, we’ve discovered several interesting properties of SMILES strings for molecules from an attention model-based standpoint, and our approaches and results may still be useful for others investigating computational chemistry.

9 Acknowledgements

We would like to thank Vineet Kosaraju, Michele Catasta, Keiran Thompson, Lillian Zhu, Gaby Li, and Amelia Woodward for their deeply helpful consultation with the project.

References

- [1] Z. Alperstein, A. Cherkasov, and J. T. Rolfe. All smiles variational autoencoder, 2019.
- [2] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [3] A. Bau, Y. Belinkov, H. Sajjad, N. Durrani, F. Dalvi, and J. Glass. Identifying and controlling important neurons in neural machine translation. *arXiv preprint arXiv:1811.01157*, 2018.

²<ftp://ftp.ncbi.nlm.nih.gov/pubchem/Compound/Extras/>

- [4] K. Clark, U. Khandelwal, O. Levy, and C. D. Manning. What does bert look at? an analysis of bert’s attention. *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, 2019.
- [5] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding, 2018.
- [6] P. Ertl. An algorithm to identify functional groups in organic molecules. *Journal of cheminformatics*, 9(1):36, 2017.
- [7] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.
- [8] V. Garcia and J. Bruna. Few-shot learning with graph neural networks. *arXiv preprint arXiv:1711.04043*, 2017.
- [9] N. Haider. Checkmol/matchmol, 2003.
- [10] W. Hamilton, Z. Ying, and J. Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pages 1024–1034, 2017.
- [11] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [12] B. Hoover, H. Strobel, and S. Gehrmann. exbert: A visual analysis tool to explore learned representations in transformers models. *ArXiv*, abs/1910.05276, 2019.
- [13] W. Hu, B. Liu, J. Gomes, M. Zitnik, P. Liang, V. Pande, and J. Leskovec. Pre-training graph neural networks. *ArXiv*, abs/1905.12265, 2019.
- [14] R. Huang, M. Xia, D.-T. Nguyen, T. Zhao, S. Sakamuru, J. Zhao, S. A. Shahane, A. Rossoshek, and A. Simeonov. Tox21challenge to build predictive models of nuclear receptor and stress response pathways as mediated by exposure to environmental chemicals and drugs. *Frontiers in Environmental Science*, 3:85, 2016.
- [15] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [16] L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han. On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265*, 2019.
- [17] S. Liu, T. Li, Z. Li, V. Srikumar, V. Pascucci, and P.-T. Bremer. Visual interrogation of attention-based models for natural language inference and machine comprehension. Technical report, Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States), 2018.
- [18] X. Liu, P. He, W. Chen, and J. Gao. Multi-task deep neural networks for natural language understanding. *arXiv preprint arXiv:1901.11504*, 2019.
- [19] T. Mikolov, M. Karafiat, L. Burget, J. Černocký, and S. Khudanpur. Recurrent neural network based language model. In *Eleventh annual conference of the international speech communication association*, 2010.
- [20] P. G. Polishchuk, T. I. Madzhidov, and A. Varnek. Estimation of the size of drug-like chemical space based on gdb-17 data. *Journal of computer-aided molecular design*, 27(8):675–679, 2013.
- [21] P. Pope, S. Kolouri, M. Rostami, C. Martin, and H. Hoffmann. Discovering molecular functional groups using graph convolutional neural networks, 12 2018.
- [22] P. E. Pope, S. Kolouri, M. Rostami, C. E. Martin, and H. Hoffmann. Explainability methods for graph convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10772–10781, 2019.
- [23] B. Ramsundar, S. Kearnes, P. Riley, D. Webster, D. Konerding, and V. Pande. Massively multitask networks for drug discovery. *arXiv:1502.02072*, 02 2015.
- [24] M. Ren, E. Triantafillou, S. Ravi, J. Snell, K. Swersky, J. B. Tenenbaum, H. Larochelle, and R. S. Zemel. Meta-learning for semi-supervised few-shot classification. In *Proceedings of 6th International Conference on Learning Representations ICLR*, 2018.
- [25] T. Rocktäschel, E. Grefenstette, K. M. Hermann, T. Kočiský, and P. Blunsom. Reasoning about entailment with neural attention. *arXiv preprint arXiv:1509.06664*, 2015.
- [26] S. G. Rohrer and K. Baumann. Maximum unbiased validation (muv) data sets for virtual screening based on pubchem bioactivity data. *Journal of chemical information and modeling*, 49(2):169–184, 2009.
- [27] A. M. Rush, S. Chopra, and J. Weston. A neural attention model for abstractive sentence summarization. *arXiv preprint arXiv:1509.00685*, 2015.
- [28] P. Schwaller, T. Laino, T. Gaudin, P. Bolgar, C. Bekas, and A. A. Lee. Molecular transformer-a model for uncertainty-calibrated chemical reaction prediction. *arXiv preprint arXiv:1811.02633*, 2018.
- [29] P. Schwaller, T. Laino, T. Gaudin, P. Bolgar, C. A. Hunter, C. Bekas, and A. A. Lee. Molecular transformer: A model for uncertainty-calibrated chemical reaction prediction. *ACS Central Science*, 5(9):1572–1583, Aug 2019.
- [30] H. Strobel, S. Gehrmann, M. Behrisch, A. Perer, H. Pfister, and A. M. Rush. Seq2seq-v is: A visual debugging tool for sequence-to-sequence models. *IEEE transactions on visualization and computer graphics*, 25(1):353–363, 2018.
- [31] A. A. Toropov, A. P. Toropova, D. V. Mukhamedzhanova, and I. Gutman. Simplified molecular input line entry system (smiles) as an alternative for constructing quantitative structure-property relationships (qspr). 2005.
- [32] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. *CoRR*, abs/1706.03762, 2017.

- [33] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [34] Z. Xiong, D. Wang, X. Liu, F. Zhong, X. Wan, X. Li, Z. Li, X. Luo, K. Chen, H. Jiang, et al. Pushing the boundaries of molecular representation for drug discovery with the graph attention mechanism. *Journal of medicinal chemistry*.
- [35] J. You, B. Liu, Z. Ying, V. Pande, and J. Leskovec. Graph convolutional policy network for goal-directed molecular graph generation. In *Advances in Neural Information Processing Systems*, pages 6410–6421, 2018.
- [36] Z. Zhou, S. Kearnes, L. Li, R. N. Zare, and P. Riley. Optimization of molecules via deep reinforcement learning. *Scientific reports*, 9(1):1–10, 2019.