

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

5-2022

An exploratory study on code attention in BERT

Rishab SHARMA

Fuxiang CHEN

Fatemeh H. FARD

David LO

Singapore Management University, davidlo@smu.edu.sg

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#)

Citation

SHARMA, Rishab; CHEN, Fuxiang; FARD, Fatemeh H.; and LO, David. An exploratory study on code attention in BERT. (2022). *Proceedings of the 30th International Conference on Program Comprehension, Virtual Event, 2022 May 16-17*. 437-448.

Available at: https://ink.library.smu.edu.sg/sis_research/7694

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.

An Exploratory Study on Code Attention in BERT

Rishab Sharma
University of British Columbia
Canada
rishab.sharma@alumni.ubc.ca

Fuxiang Chen
University of British Columbia
Canada
fuxiang.chen@ubc.ca

Fatemeh Fard
University of British Columbia
Canada
fatemeh.fard@ubc.ca

David Lo
Singapore Management University
Singapore
davidlo@smu.edu.sg

ABSTRACT

Many recent models in software engineering introduced deep neural models based on the Transformer architecture or use transformer-based Pre-trained Language Models (PLM) trained on code. Although these models achieve the state of the arts results in many downstream tasks such as code summarization and bug detection, they are based on Transformer and PLM, which are mainly studied in the Natural Language Processing (NLP) field. The current studies rely on the reasoning and practices from NLP for these models in code, despite the differences between natural languages and programming languages. There is also limited literature on explaining how code is modeled.

Here, we investigate the attention behavior of PLM on code and compare it with natural language. We pre-trained BERT, a Transformer based PLM, on code and explored what kind of information it learns, both semantic and syntactic. We run several experiments to analyze the attention values of code constructs on each other and what BERT learns in each layer. Our analyses show that BERT pays more attention to syntactic entities, specifically *identifiers* and *separators*, in contrast to the most attended token [CLS] in NLP. This observation motivated us to leverage *identifiers* to represent the code sequence instead of the [CLS] token when used for code clone detection. Our results show that employing embeddings from identifiers increases the performance of BERT by 605% and 4% F1-score in its lower layers and the upper layers, respectively. When identifiers' embeddings are used in CodeBERT, a code-based PLM, the performance is improved by 21–24% in the F1-score of clone detection. The findings can benefit the research community by using code-specific representations instead of applying the common embeddings used in NLP, and open new directions for developing smaller models with similar performance.

CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence.**

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICPC '22, May 16–17, 2022, Virtual Event, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

KEYWORDS

pre-trained language models, BERT, CodeBERT, attention

ACM Reference Format:

Rishab Sharma, Fuxiang Chen, Fatemeh Fard, and David Lo. 2022. An Exploratory Study on Code Attention in BERT. In *30th International Conference on Program Comprehension (ICPC '22)*, May 16–17, 2022, Virtual Event, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Pre-trained language models (PLMs) such as BERT [10] and RoBERTa [27] are based on Transformer networks [43] and have been widely used in the field of Natural Language Processing (NLP). These language models are trained on a large corpus of data and are able to capture the general understanding of a language [33]. More recently, the software engineering community has employed language modeling on tasks such as code retrieval, program repair, code documentation, code clone detection, and bug detection [12, 13, 19, 20]. The results obtained by fine-tuning the language models pre-trained on code corpora outperformed the previous state-of-the-art models [12, 14, 28]. There have been attempts [12, 14] to develop code-specific PLMs, which can leverage existing techniques in Software Engineering (SE) to enrich the underlying Transformer networks for SE-related tasks. However, all these works are based on studies in natural language and they do not investigate how code is modeled differently.

Under the hood, these transformer-based language models use the attention mechanism, which enables the model to attend to more important tokens in a sequence. This is used in multiple software engineering studies for developing models or explaining the models' predictions [3, 15, 23, 25, 38]. The importance of this architecture with attention mechanism can be derived from the fact that employing the vanilla Transformer based networks without any additional techniques like utilizing abstract syntax tree and information retrieval can generate state-of-the-art results in various tasks such as code summarization, code retrieval, code clone detection, and program repair [1, 2, 12, 14, 46]. However, there is no study to explore how the PLMs attend toward different tokens in code and whether the assumptions made in NLP for representing a text sequence apply to code, despite the differences between programming languages and natural languages. For example, code comprises of different constructs such as identifiers and data-types, and how the PLMs model these constructs is unknown.

Although there are studies in software engineering [7] that show the similarities and dissimilarities between the programming languages and natural languages, assuming similar behavior in code by directly applying the NLP knowledge may not be appropriate; and it is necessary to study how the Transformer architecture using language models learns the semantic and syntactic information in code. The literature on this topic is limited. It is only recently that Karmakar and Robbes [21] studied the comprehension abilities of four different pre-trained models (BERT [10], CodeBERTa [48], CodeBERT [12] and GraphCodeBERT [14]) trained on code. They compare the performance of these models from the perspective of probing classifiers without investigating the details of what the architecture learns about code.

This study evaluates the BERT model, trained on code corpus. BERT is chosen as it is the basis of RoBERTa used by CodeBERT and other PLMs developed for code. In addition, we want to compare our code-specific findings with what BERT learns about natural language in the study of Clark et al. [8]. We evaluate the BERT model that we pre-train on code as a classical NLP pipeline. An NLP pipeline analyzes the language based on its linguistic features, which can contain semantic and structural properties [30, 41]. In our study, we explore the **attention behavior** of BERT on different syntactic types (i.e., code constructs). We want to understand **which layers** the model learns the semantic and syntactic knowledge of code. To compare our findings (BERT trained on programming languages) with (BERT trained on natural languages), a direct comparison is made with an existing NLP study by Clark et al. [8]. Our study uses BERT as it is a popular and strong baseline used in NLP and Software Engineering (SE) areas. Moreover, other variations of BERT such as RoBERTa [27] are based on BERT and they have a similar architecture that is used in many PLMs on code [12, 14]. Our results show that, unlike NLP, the [CLS] and [SEP] tokens are not attended much in code. This is important as the learned embedding of [CLS] and [SEP] are frequently used in the classification tasks in both fields of NLP and SE. However, we observed that the [CLS] is not the best token to capture the information of the code sequence for classification tasks. On the other hand, BERT attends more toward *identifiers* (e.g., class names) and *separators* (e.g., `;`, `{`) in code. This led us to the hypothesis that the embeddings learned for identifiers in code can be a better representation in classification tasks related to code. To validate our hypothesis, we conducted experiments on code clone detection, where we compare the results obtained using the embeddings from [CLS] and *identifiers*. When the identifier’s embedding is used, the performance is improved by 605% and 4% in F1-score, using embeddings from the lower layers and the upper layers of BERT, respectively. We apply our findings on CodeBERT [12], a Transformer based PLM developed for code representation, which uses a different pre-training objective. Similar behavior is observed when using the identifier’s embedding, and the results of code clone detection are improved by 21–24% in F1-score compared to when the [CLS] embedding is used.

The contributions of this paper are listed as follows:

- We present a first study that explores the **attention comprehension** abilities of BERT pre-trained on code.
- We compare BERT’s learning in the modeling of natural language *English* and the programming language *Java*. Our

findings reveal some similarities and differences in modeling the languages by the PLMs.

- We study whether BERT can intrinsically learn code structure. Our analysis shows that the models can understand the general syntactic structure by predicting correct syntactic entities; however, they cannot correctly predict the valid tokens with high accuracy.
- We propose to use the embeddings generated from identifiers in place of the vanilla embeddings used in BERT. When employed for code clone detection, these embeddings create better results than the vanilla techniques used in the PLM.

We have open-sourced our experiments for replication of the results and further usage in the software engineering community¹.

Significance of Study: Our study provides new insights on two important entities in code: *identifiers* and *separators*. This also includes how they are modeled in PLMs. The representation of these code constructs can be used to generate better models for code clone detection. The generated embeddings are effective as they increase the performance significantly using the lower layers of BERT. Researchers can use the findings presented in the paper to use the current PLMs and obtain better results. Moreover, though this is not the focus of our research, these findings can be used to reduce model sizes, addressing the lack of computational accessibility and enable a wider audience to use PLMs in practice, e.g. by having 6 layer models with similar performance to 12 layer models.

The rest of this paper is organized as follows. In section 2 we detail the necessary background. Study design and the experimental setup are presented in section 3, followed by answers to each research question in sections 4, 5, and 6. Related works are discussed in section 7, followed by threats to validity in section 8. We conclude the paper in section 9.

2 TRANSFORMER AND BERT ARCHITECTURE

In this section, we explain the necessary background.

2.1 Transformer

Transformer forms the fundamental units of BERT, a popular PLM used in many studies [10]. A Transformer in a sequence generation setting uses multiple stacks of Encoder and Decoder [43]. Transformers primarily leverage Multi-Head Self Attention and Feed-forward neural network to model the relationship within the text [43]. BERT solely uses the Transformer’s encoder stack due to its pre-training objectives used in its pre-training [10].

Multi-Head Self Attention Layer: The Transformer encoders use the self-attention [43] mechanism to represent the input sequence by relating every input sequence element with each other. For the given input sequence S ($S = T_1, T_2, \dots, T_n$), of length n , each input token T_i is first converted into its vector embedding x_i .

Each encoder layer consists of three sets of weights matrices, $W(q)$, $W(k)$, and $W(v)$ that is used for query, key, and value, respectively. Initially, these matrices are randomly initialized and they are trainable parameters. The vector multiplication between input embedding x_i and $W(q)$, $W(k)$, and $W(v)$ matrices generates

¹<https://github.com/fardfh-lab/Code-Attention-BERT>

unique query $q(i)$, key $k(i)$, and value $v(i)$ vectors, respectively.

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

The attention scores are obtained by performing dot product between vector $q(i)$ and vector $k(i)$. This score is then divided by the square root of the dimension of the input embedding d_k . The softmax function is applied to the attention scores to obtain attention distribution as shown in (1). Each softmax score is multiplied with the value vector $v(i)$ to generate weighted vectors added to get the final vector $z(i)$. This same task is repeated for each token multiple times, to form multi-headed attention. Each attention head learns a different vector space which is then combined by concatenating the information learned by each attention head.

Position Embeddings: Languages are sequential, and changing the order of words may lead to a completely different meaning or an erroneous sentence, which does not have any congruous meaning. In Transformer, a complete sequence of text forms the input into the multi-head attention layer, which helps in promoting parallelization. To maintain and learn the syntactic correctness of the language, an additional type of embeddings, position embeddings, is introduced. Specifically, the word embedding x_i and its corresponding position embeddings p_i are added, $xp_i = x_i + p_i$. Embedding xp_i is then further fed into the multi-head attention layer:

Residual Connection and Add Normalize: The output of multi-head attention is added to the residual connection from $x(i)$, and layer normalization is performed on the added vector.

Feed Forward Neural Network Layer: The output from the Add Normalize layer is then fed into the feed-forward neural network, which has a common architecture for all the tokens. However, this step is done in parallel on each token as shown in equation (2).

$$\text{FFN}(x_i) = \max(0, x_i W_1 + b_1) W_2 + b_2 \quad (2)$$

The output is then added and normalized. This finally completes a single feed-forward propagation in one encoder unit. The output from this encoder unit is fed into the next unit, and the same processing is done on it.

2.2 BERT: Bidirectional Encoder Representations From Transformers

BERT is a PLM trained with the core motivation to build a bi-directional model, which can learn the general understanding of the language and then it can be fine-tuned for various downstream tasks.

BERT-Pretraining: BERT is pre-trained on two unsupervised tasks (i) Masked Language Modelling (MLM) and (ii) Next Sentence Prediction (NSP). Originally, BooksCorpus and English Wikipedia datasets were used to train BERT [10]. The datasets contain books and Wikipedia information that is available online.

Masked Language Modeling (MLM): Unlike other neural language models, BERT is pre-trained using Masked Language Modeling. In masked language modeling, the task is to predict a randomly masked token in a sequence rather than predicting the next token in sequence. Within Masked Language Modeling in BERT, 15% of the tokens are randomly masked. Masking of the tokens follows a probability pattern where a token is either replaced with (1) a [MASK] token 80% of the time, (2) a random token 10% of the time,

or (3) it will remain unchanged for 10% of the time.

Next Sentence Prediction (NSP): BERT is also pre-trained on a binary next sentence prediction task. This task focuses on training a language model that can discern whether the given sentence is the next sentence of the current sentence. This is done by randomly replacing the next sentence 50% of the time with random sentences. This would enable the model to discern the difference between the natural coherence of the sentences, which motivates the model to learn the general understanding of a language.

Special Tokens in BERT: BERT introduces special tokens namely [CLS] and [SEP] tokens as its input formatting. [CLS] is appended at the start of the sentence, which is supposed to learn knowledge of the whole sentence. [SEP] token is used as a “separator” between sentences to highlight the end of a sentence. The representation from the [CLS] token is also used during the NSP pre-training task of the BERT model, as a “classification” token to predict whether a sentence is the next sentence of a previous sentence.

3 BERT ATTENTION FOR CODE

The Transformer-based model primarily uses the multi-headed attention to model languages [10, 27]. Therefore, attention plays a vital role in training the PLMs, and thus it is important to study the attention behavior of PLMs on the given text information. An attention distribution for a given sentence can quantify the Transformer model’s importance to a particular token. A higher attention value on a specific token means that token has a higher weight in determining the representation of a given sentence, making it a significant token to learn and represent a sentence. Therefore, knowing the important tokens within the code is important for extracting meaningful information from the PLMs. To understand and visualize this attention information, we plot the attention maps for different kinds of token types, similar to a study released for the English language by Clark et al. [8]. We use this study as an initial comparison of Transformer learning between code and natural language, and design different experiments to understand the behavior of this architecture for programming languages.

3.1 Experimental Setup

There are two variants of BERT: BERT-Large and BERT-Base [10]. BERT-Base model contains 12 stacked encoders, with 12 attention heads in each layer. In our study, we use the BERT-Base cased language model as a direct comparison against the natural language, English. The cased language model is used so that the BERT model can learn the Java coding conventions for the identifiers, which promotes the use of different casing styles. As we intend to compare our results of the BERT model trained on code with natural language, we need to keep the models the same. There are multiple implementations available for BERT and its variants trained on code, which are pre-trained with different objectives (e.g., MLM and Replaced Token Detection).

However, for the comparison, we required a BERT model that is pre-trained with the same objectives as the one used for the NLP study [8] and also should purely be pre-trained on code. We are unable to find an open-source pre-trained BERT-Base model that is trained solely on code. Therefore, we trained a BERT-Base cased model from scratch. All the pre-training decisions like the number

of layers in the encoder and decoder, the casing of tokens, and the pre-training objectives are consistent with the experimental settings in the NLP study by Clark et al. [8]. For pre-training the BERT model, we used the benchmark dataset released by Husain et al. [18]. This dataset contains a large corpus for six different programming languages collected from open source repositories available on GitHub. However, we only confine our work to the Java programming language in this study (we do not require a multilingual model) to keep it consistent with the NLP study where the BERT model is only trained on a single natural language. Java is chosen as the language in our study because it is a widely adopted programming language to study program comprehension in the software engineering community. It is also a commonly used programming language in the industry. The statistics of the dataset used is shown in Table 1.

Table 1: Dataset details for pre-training BERT

Split	# Records in the Dataset
Train	164,923
Validation	5,183
Test	10,955

For training the BERT-Base based model, we used the open-source implementation of BERT that is available from its authors, Devlin et al. [10]. In this study, as we focus only on code representation using BERT, we remove the JavaDoc comments, the inline comments, and the block comments from the dataset. This cleaned dataset only contains code and it is used for training the BERT model. Further, to make the dataset eligible for the next sentence prediction task, we keep the separation of the code sentences in different lines, as mentioned on the BERT’s official implementation repository² and this separation process is also used in previous study [26]. We then train BERT using this dataset. The training continues for four consecutive days on a Tesla V100 GPU with 32GB memory. The performance of the BERT model on the validation dataset for the two pre-training objectives, Masked Language Modeling and Next Sentence Prediction, is shown in Table 2. The results obtained for BERT’s pre-training for the NSP objective on code is 94.75%, which is close to the original pre-training result for BERT model on English, as reported by the authors of the BERT model. This result confirms that the model we have trained is consistent with the original BERT model on English language. The accuracy achieved by BERT for MLM is not reported in the original paper. Therefore, we cannot make any direct comparison for this pre-training objective. Nonetheless, we believe that 87.44% is a reasonable accuracy for predicting the masked tokens, as it involves the correct prediction of the code token in a sentence.

In the remaining paper, the BERT model that we refer to is the model that we have pre-trained on Java.

3.2 Research Questions

This study investigates answers to the following research questions, which require analysis of the BERT’s attention behavior and comprehension abilities for code.

²<https://github.com/google-research/bert>

Table 2: Accuracy of pre-trained BERT-Base on Valid dataset. The results are close to the original pre-training results reported in BERT

Pre-Training Task	% Accuracy
Masked Language Modeling	87.44
Next Sentence Prediction	94.75

RQ1 (Attention Behavior): What is the surface level behavior of BERT attention over code? In this research question, we study the attention scores set by BERT on the special tokens available within BERT. Moreover, we also study the divergence of information learned by different attention heads in various layers of BERT. The relative position on tokens is used to investigate the short or long-range attention over different layers.

RQ2 (Code Construct Relationships): How does BERT encode the code-specific relationships among different code constructs? In this research question, we explore the code constructs that receive the highest attention, which can reveal important code constructs learned by the model. Furthermore, we investigate how each code construct, specifically identifiers, distributes its relationship with other code artifacts like separators and keywords in different layers. Moreover, we study whether BERT can encode syntactic information of code by using a non-parametric probing.

RQ3 (Semantic Knowledge): How does BERT learn the semantic knowledge in code? In this research question, we investigate which layers within the BERT model can learn the semantic understanding of code. We do this by analyzing the performance of each layer for the code clone detection task. Finally, based on our findings from RQ1 and RQ2, we leverage the embeddings learned from identifiers to better represent code in place of the vanilla [CLS] token.

3.3 Generating Attention Maps

Once BERT is pre-trained on the Java dataset, we use the model’s trained weights to extract the attention distribution. For extracting the attention scores, following Clark et al. [8], we do not use the train dataset for understanding the attention mechanism. Instead, we select samples from the test dataset. As BERT uses sentence piece tokenization of text, it can lead to sub-tokenization of a token into multiple tokens depending on the vocabulary used within the tokenizer. Therefore, we can sample code examples using two different strategies (i) Select only code examples from the test dataset for which the length before and after the tokenization remains the same, (ii) Randomly sample code where tokens are sub-tokenized and are later combined for attention analysis. We conducted the same experiments using both strategies; however, the results are similar irrespective of the sampling strategy. Therefore, to avoid redundancy, we only include the experiments of the second strategy and combine the tokens for the analysis.

4 ATTENTION BEHAVIOR

Here, we would only focus on the surface-level behavior of BERT’s attention on code. We follow the surface level analysis as conducted

in [8], and only include the attention over special tokens, the relative positioning of the tokens, and the similarity and difference in attention distribution over different layers and heads.

4.1 Attention on Special Tokens

Before studying the attention behavior of code constructs, we look at the attention behavior over the special tokens introduced by the BERT's pre-processing. It is important to study these special tokens as they play an important role in pre-training the BERT model. Two special tokens are introduced to feed input to the BERT model, namely [CLS] and SEP. The [CLS] token's representation is used during pre-training for the next sentence prediction objective and assumed to contain the complete information about the sequence, and the [SEP] token acts as a delimiter between two sentences.

To visualize the attention distribution of these tokens, we apply the following process. The BERT architecture we use has 12 layers, and each layer contains 12 attention heads. As discussed previously, each attention head is trained to learn different aspects of the input, and when all the heads are combined, they should reflect the complete knowledge. Therefore, for each layer and each head in the layer, we aggregate the attention scores of the model for the [CLS] and [SEP] tokens, separately. We aggregate the total attention scores on these special tokens (the scores for each of the [CLS] and [SEP] are computed separately), which is the attention given to these tokens by all the other words in the sequence. This score is then normalized with the token's total number of occurrences, giving us the average attention put on each of [CLS] and [SEP] tokens in each given code. The results are shown in Figure 1. The multi-scatter plot on the y-axis of each layer shows the values of different attention heads at each layer.




Figure 1: Attention on special tokens [CLS] and [SEP]

As shown in Figure 1, both of these special tokens have very low attention values, which means BERT has little focus on the [CLS] and [SEP] tokens when modeling code. This reduces the significance of special tokens in learning the relationship of code tokens. This result is contrary to what has been seen in the NLP study [8]. In NLP, the average attention on [CLS] and [SEP] ranges between 0 to 0.6 over different layers. This score ranges only between 0 to 0.10 for code.

For the [SEP] token, the average attention increases between Layers-4 to Layers-6, and then there is a consistent decline until layer 10, and then remains uniform for the remaining layers. Overall, the [SEP] token consistently gets more attention over the [CLS] token in all layers. However, in the case of NLP, Clark et al. [8] reported that the initial layers until Layer-4 attend more towards [CLS] over the [SEP] token, which is not the case as seen in Figure 1. We see a gradual decrease for the [CLS] token until Layer-5 and then a gradual increase until Layer-9.

4.2 Attention on Relative Positioning of Tokens

The attention score of the relative positioning of the tokens can reveal the range (long and short) of attention distribution learned by the model. Attention towards immediate neighbor tokens means the model primarily understands the local context of the code. In this study, we find the model's attention that is put on the current token, previous token, and the next token. As shown in Figure 2, BERT focuses less towards the token itself while learning its representation when compared to attention towards its left and right token. In the initial layers, the concentration on the token itself is lower, which gradually increases in the middle layers (Layer-5 to Layer-9). At Layer-11, the attention of the current token increases and it is close to the attention value of its left and right token.




Figure 2: Attention on relative position of tokens

On average, 30% of attention is put by BERT on the left and right token. Uniform attention values are seen over each layer. Also, it demonstrates that the BERT for code model attends better locally than the long-range attention. This is similar to what we see in NLP, as BERT attends more towards local tokens in the English language.

4.3 Attention Redundancy at Attention Heads and Layers

BERT uses a multi-head self-attention mechanism to learn the representation of tokens. Multiple attention heads within the same encoder layer are expected to learn different language features. The variety of information learned by these attention heads is then concatenated to generate the representation for the language.

Hence, it is important to study how much knowledge these attention heads learn from each other.




Figure 3: Attention redundancy at attention heads and layers

To understand this information, the Jensen-Shannon divergence [31] between different attention heads in the BERT model is calculated. Jensen-Shannon divergence finds the similarity between two probability distributions. Therefore, we use the probability distribution of these attention heads to learn the differences. The more the distance between the two heads, the more disparate information is retained by them. In total, BERT contains 144 attention heads. Therefore, we have a 144×144 matrix to represent such information, as shown in Figure 3. The middle layers retain diverse information compared to the lower and higher layers. Similar to the NLP study, redundancy in the attention distribution insinuates that not all attention heads learn unique information. Therefore, there is scope for applying model distillation strategies that can help reduce the model size and maintain equivalent performance. The only difference we see between Java and English is the absolute divergence values. The divergence in the information learned in different attention heads in English is slightly more than the code, possibly due to the repetition of code tokens and code syntax [7].

RQ1 Summary: We studied the surface level behavior of BERT for code. Our experiments show that this behavior for code is mostly similar to NLP. Therefore, surface-level behavior is dependent more on the architecture of BERT, and the language differences have a minor effect on the surface-level behavior. The primary finding is the difference between code and NLP in the amount of attention learned by BERT for the special tokens [CLS] and [SEP], which is much less for code; indicating that the [CLS] representation cannot efficiently capture the information about the code sequence.

5 CODE CONSTRUCT RELATIONSHIPS

Code has rich grammar and syntax. It contains multiple syntactic constructs such as *identifiers*, *separators*, and *data types*. These code constructs are not available in natural languages. Therefore, we study the comprehension abilities of BERT for these syntactic types and the relationship of each syntactic entity with other code constructs.




Figure 4: Attention on identifier (IDF) and separators SEPS

5.1 Analysing Attention on Code Constructs in Code

Syntax plays an important role to increase the performance of neural models on different software engineering tasks, including code comment generation [17, 24, 50], code clone detection [12, 14], program repair [47], software vulnerability detection [16], type prediction [29], and bug detection [34]. Here, we study the attention scores of BERT on the code constructs.

We draw an attention distribution plot similar to the plot for special tokens, [SEP] and [CLS] as done in section 4.1. In this plot, we collect the attention distribution for each of the syntactic types (i.e., code constructs). In order to tag each token to a syntactic type, we use the JavaLang Parser³. The JavaLang parser provides both parser and tokenizer for the Java programming language and it has been used widely in other software engineering studies [26, 45, 50]. The syntactic types included by the parser are *identifiers*, *separators* (e.g., ‘,’; ‘;’; ‘{’; ‘}’), *operators* (e.g., =, +, -), *data-types* (int, float, double, etc.), *keywords* (e.g., abstract, continue, const), and *access-modifiers* (e.g., public, private). We group all the *separators* into a single token SEPS. Note that the [SEP] token in Figure 1 is different from SEPS in Figure 4. The former is a special token introduced by BERT to act as a delimiter between two sentences, and the latter is for the syntactic types (e.g., ‘,’; ‘;’; ‘{’; ‘}’) available within code.

The plots are shown in Figures 4 and 5. Figure 4 demonstrates that identifiers (IDF) and separators (SEPS) get the most amount of attention. The higher attention values over identifiers and separators demonstrate that BERT considers them as important tokens for code modeling. Therefore, both of these tokens should encode the latent knowledge of code better than the [CLS] and [SEP] special tokens.

Interestingly, BERT models most of its attention over identifiers inherently, which could be a possible reason for the success of the BERT model. Between identifiers and separators, identifiers have higher attention values ranging between 30%-45%, until Layer-11. At the final layer, the separators tend to get the most amount of attention. A possible reason for high attention values for identifiers and separators could be that both types are the most frequently

³<https://github.com/c2nes/javalang>

occurring tokens in the code. Also, the same token repeats at different positions with the same code sample due to the structure of the code [7].

Literature on program comprehension has also acknowledged the importance of identifiers for easier program understanding [4, 11, 32, 37]. Based on the obtained results of Figure 4, a question is that “when modeling code, can using the IDF embeddings boost the performance of the same model over using the embeddings from [CLS]”? We answer this question in section 6.2.




Figure 5: Attention on Syntactic types including Operators (OP), Data-type (DTP), Keywords (KEY) and Modifier (MOD)

The attention plots of other syntactic tokens, including operator (OP), data-type (DTP), keywords (KEY), and access-modifiers (MOD) are shown separately in Figure 5. They are shown separately in Figure 5, as their attention scores are very low compared to identifiers and separators. Data type gets the least attention, followed by modifier, keyword, and operator. Such low attention values for these code constructs are surprising because Java is a statically typed language, and knowing data types is important for the correct compilation of code. Moreover, all of these syntactic types are related to the structural part of the code.

5.2 Modeling the Relationship Among Syntactic Tokens and Identifiers

The results shown in the previous section indicate that BERT has the most attention on identifiers. As identifiers have previously shown to be important code constructs for program comprehension [4, 11, 32, 37], we analyze where identifiers focus when BERT models the identifiers’ representation.

To analyse this attention distribution, we extracted the attention matrix of indexes where identifiers were available in code. Then, we leverage the identifier attention distribution, to model its relationship with other identifiers (IDF), separators (SEPS), operators (OP), data type (DTP), keywords (KEY) and modifiers (MOD) as shown in Figure 6.

As seen in the top plot of Figure 6, from Layer-1 to Layer-11 identifiers attends most towards identifiers, similar to the overall attention behavior of the BERT model. This is followed by separators and operators within the code. Again, data-type, keyword, and modifiers have the least attention values. As discussed in section 4.2, the attention is generally focused on local context. Hence, this attention behavior of identifiers should also be from the local context.




Figure 6: Modeling the relationship among syntactic tokens with identifier



Figure 7: Attention of identifiers observable and t at layer 10. This representation has been generated using Bertviz [44]

Figure 7 shows an example of a code sequence and the attention distribution of tokens. There are two identifiers `observable` and `t`, for which the attention distribution over other tokens is highlighted, with links to those tokens. The darker the link, the higher is the attention value towards that token. We observe that these identifiers attend mostly on themselves. Each of them models its relationship with other local tokens (i.e., its immediate left and right tokens) to learn their representation, focusing on longer distance tokens with low attention values. Identifiers are generally associated with a data type; however, BERT fails to understand such associations intrinsically as data-type gets the least attention.

5.3 Predicting Syntactic Entities Within Code

To further investigate the abilities of BERT to capture the syntactic information of code, we conducted an experiment to understand whether BERT can generate the correct syntax of Java programming language that is compiled using a compiler. Please note that here we do not want to find the accurate token; rather, we aim to know whether BERT can accurately predict the syntax type of the missing tokens that follow the Java language’s syntactic rules. We use a non-parametric probing task as used in [12]. In this task, we mask the tokens available in the code. However, we mask these tokens based on their syntax type. For instance, given a code sample, we first mask all the *identifiers* within the code and test the ability of the BERT model in predicting the identifiers in the masked spaces. Similarly, we repeat this masking for *data-type*, *decimal integer*, *keywords*, *modifier*, *operator* and *separators*. The results are shown in Table 3.

Table 3: Predictions for Syntactic Entities

Type	Precision	Recall	F1
data type	0.66	0.96	0.78
decimalinteger	0.59	0.69	0.61
identifier	0.98	0.88	0.93
keyword	0.68	0.95	0.79
modifier	1.00	0.99	0.99
operator	0.93	0.98	0.96
separator	0.98	0.97	0.98

We evaluate the performance of this task using Precision, Recall, and F1 scores as follows: Precision = $TP/(TP + FP)$; Recall = $TP/(TP + FN)$; F1 = $2 \times \text{Precision} \times \text{Recall} / (\text{Precision} + \text{Recall})$. Here, True Positive (TP) is the number of tokens correctly identified to replace the MASK tokens. False Positive (FP) is the number of samples incorrectly labeled as belonging to a group. False Negative (FN) is the sample model incorrectly marked as not belonging to a group.

Overall, BERT can correctly predict token types for identifier, separator, modifier, and operators. Syntactic types decimal integer, data-type, and keyword have lower performance than other types. The lower scores of these types can be related to the lower attention values they receive during the modeling as described in Section 5.1. The obtained scores for modifiers are high, despite the lower attention values. This can be related to the fact that the total number of modifiers in Java is less than five, making it an easy prediction.

RQ2 Summary: BERT can learn different code constructs with varying performance. Mostly, BERT focuses on *identifiers* and *separators*. Therefore, the representation of identifiers can be useful for the downstream tasks. Also, the representation learned for an identifier is mostly contributed by the token itself or its immediate left and right tokens.

6 SEMANTIC KNOWLEDGE

Semantic learning is pivotal to any language. It helps the model learn representation meanings/ relationships within the learned space. In this section, we use code clone detection, a classification task that requires semantic information of code, to learn semantic information captured by BERT in each layer. Code clones is chosen as it provides the ability to study both the semantic and structural

properties of code using one task. In this experiment, we use the embedding of the code sequence at each layer of BERT for identifying code clones. Code clone detection is chosen as it requires the model to capture some semantics about the code, where code samples can have structural or semantic equivalences.

This experiment is two fold: First, we assess the ability of BERT in capturing the semantics of code in different layers, using the learned embedding of the [CLS] token. In the second part, we seek the answer to our question: whether we can improve the performance by using the embeddings from IDF instead of [CLS], based on our findings in RQ1-2, which are explained below.

6.1 Semantic Representation of Code for Code Clone Detection Using [CLS]

Code clones are code samples that are identical to each other [6, 22]. These code samples can have structural or semantic equivalences [40]. Generally, code clone detection finds target code that has similar attributes to the given code [40]. If the code samples have the same functionality, they are tagged as semantic clones. Otherwise, they are classified as non-clones of each other. For the code clone detection, we use the pipeline made available by Lu et al. [28], with few changes. The original pipeline was built for ROBERTa [27], which is changed to accommodate the BERT model. The pipeline uses the BERT model to initialize the encoder, and then a linear head is used over the pre-trained model for the classification into clones and non-clones. The classification model is trained end-to-end, and, similar to works in NLP [41, 42], we freeze the weights of the BERT model because we want to learn the model’s capabilities on the target task without further training the BERT model. Therefore, the weights of the encoder are **not** fine-tuned on the task, and we can only extract the embeddings to assess the semantic capabilities of the pre-trained model. The linear head over BERT encoder is still fine-tuned on the task of code clone detection. The linear head is not originally part of the BERT model and it is only introduced as a projection for BERT representation. The linear head uses the hidden representation from BERT to detect the code clones. The fine-tuning is done for three epochs, following similarly to how fine-tuning was done in the original BERT paper [10]. For fine-tuning, we use the Big-Clone-Bench dataset [40]. The statistics of the dataset is shown in Table 4.

Table 4: Dataset details for Code Clone Detection

Split	# Records in the Dataset
Train	901,028
Validation	415,416
Test	415,416

To understand the semantic comprehension abilities of each layer, we extract the output hidden representation of the [CLS] token starting from Layer-1 until Layer-12. As explained previously, the embedding of [CLS] token captures the information about the input sequence and is used in classification tasks. The linear head then uses the extracted hidden state over the BERT encoder. The fine-tuning results showing the predictions of BERT for code clones are shown in Figure 8.

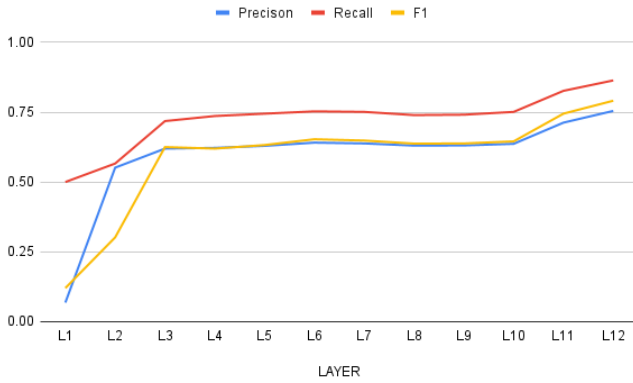


Figure 8: Performance of BERT for Code Clone Detection with [CLS] token embedding

Table 5: Using BERT *identifier* embedding for code clone detection

Layer	Precision	Recall	F1
Layer1	0.8136 (1090%)	0.9013 (80%)	0.8483 (605%)
Layer2	0.8243 (49%)	0.9150 (62%)	0.8603 (185%)
Layer3	0.8264 (33%)	0.9173 (28%)	0.8625 (38%)
Layer4	0.8333 (34%)	0.9170 (25%)	0.8674 (40%)
Layer5	0.8331 (32%)	0.9211 (24%)	0.8686 (37%)
Layer6	0.8307 (30%)	0.9211 (22%)	0.8668 (23%)
Layer7	0.8259 (29%)	0.9201 (23%)	0.8629 (33%)
Layer8	0.8266 (31%)	0.9198 (24%)	0.8634 (35%)
Layer9	0.8243 (31%)	0.9182 (24%)	0.8613 (35%)
Layer10	0.8293 (30%)	0.9182 (22%)	0.8649 (34%)
Layer11	0.8207 (15%)	0.9201 (11%)	0.8591 (15%)
Layer12	0.8365 (11%)	0.9240 (7%)	0.8719 (10%)
Layer-Pool	0.8405 (4%)	0.9260 (3%)	0.8754 (4%)

We evaluate the performance of the code clone detection using evaluation metrics for classification tasks: Precision, Recall, and F1 score, as defined in previous sections.

As seen in Figure 8, the lower layers have lower performance compared to the upper-level layers. This means that upper layers are better at representing the semantic relationship of the code. Therefore, semantic tasks should use upper-level layers to represent code to achieve the best performance. Interestingly, the performance remains uniform in the middle layers (4-10). This can be related to the uniform attention of identifiers, values we see in the Figure 4.

Table 6: Using CodeBERT *identifier* embedding for code clone detection

Layer	Precision	Recall	F1
CodeBERT _[CLS]	0.6337	0.7632	0.6294
CodeBERT _{IDF}	0.7263	0.8631	0.7619
CodeBERT-MLM _[CLS]	0.5101	0.8369	0.6339
CodeBERT-MLM _{IDF}	0.6887	0.9319	0.7921

6.2 Semantic Representation of Code for Code Clone Detection Using Identifiers

Current works [12, 14, 26] use the representation of the special tokens, i.e. [CLS] for code representation in downstream tasks for further fine-tuning. However, based on the attention distribution we saw in sections 4 and 5, we find that most of the BERT’s attention is on the *identifiers* in code, and the attention on [CLS] token is very low. In this part, we evaluate representing code using the identifiers’ embeddings. We perform an experiment similar to the one in Section 6.1 on the performance of code clone detection for each layer, instead, using the IDF representation for code. To incorporate the representation of identifiers, we calculate a weighted sum of all the identifiers present within the code. Each identifier embedding is normalized with its attention value and then added to form the weighted attention embedding of the identifiers. Finally, we extract the weighted embedding of the identifier from each layer and use it inside the classification head. Table 5 represents the results for each layer. The number inside parenthesis shows the performance change percentage compared to the model that uses representations of the [CLS] token. The results show an increase in the performance for all metrics at each layer. Using the IDF representation increases the scores for the initial layers significantly, making the difference in the results of the lower and upper layers very small. Comparing this to what we saw in Figure 8, the Precision and Recall scores are below 75 for Layer-1 to Layer-11. The performance of the pooled layer is also increased by 3–4% for all scores.

To evaluate the benefits of using embeddings from IDF over [CLS], we repeat our experiments on CodeBERT [12]. CodeBERT is chosen as it is a pre-trained model on code and it is previously evaluated on the same task by its authors, using the [CLS] token. The experiments are conducted on both variations of CodeBERT, namely, CodeBERT-MLM and CodeBERT (which considers two training objectives) [12]. We use the CodeBERT models available on HuggingFace and use the same code made available by the authors for fine-tuning on code clone detection. Similar to the pre-processing we applied for BERT to pre-trained it on code, we tokenize the code first using the JavaLang tokenizer. The tokenized code is used for code clone detection. This tokenization is necessary to extract and map the tokens to their respective syntactic types. We experimented only for the pooled layer of CodeBERT as it was the best performing model in our experiments conducted on BERT. The results are shown in Table 6, where the subscript to the models’ name indicates the representation of the token that is used. When embeddings from identifiers are used, the performance of CodeBERT is increased for both variations of CodeBERT. CodeBERT results are improved by 14.61%, 13.08%, and 21.05% for Precision, Recall, and F1 scores, respectively. The three metrics in order are 35.01%, 11.035%, and 24.95% for CodeBERT-MLM. Note that CodeBERT is based on RoBERTa [27], which is a Transformer model based on BERT.

RQ3 Summary: BERT can learn the semantic characteristics of code. The lower layers contain semantic knowledge, whereas the upper layers in BERT are rich in semantic context. Employing embeddings generated from *identifiers* in place of [CLS] token for the code clone detection task can boost the performance of the model by 4% in the pooled layer and achieve results in the initial

layers (0.848 F1 score), which are close to the final results (0.875 F1 score). This finding is not bound to BERT and improves the performance of CodeBERT (by 21–24% F1 score), which is based on RoBERTa when using IDF embeddings over [CLS].

7 RELATED WORK

Attention-based neural networks are extensively used in software engineering [3, 23]. Attention networks within the deep learning models help the model focus more on the important tokens, which can eventually affect the predictions made by the model. An attention mechanism would assign higher attention values/energy to important entities inside the input and its analysis or value in different studies to build a model or explain its predictions. LeClair et al. [23] provide attention heat maps for the predictions of their model for code summarization to explain the prediction based on the amount of attention the model puts on relevant tokens. Alon et al. [3] use an attention framework to select import Abstract Syntax Tree paths between the node to predict the appropriate function names. Similarly, Haque et al. [15] leverage the attention framework to select relevant parts of the context file to generate summaries. Shuai et al. [38] learn interdependent representations for code and query with an attention mechanism for code search, and Li et al. [25] leverage attention-based network for bug detection. Unlike previous works, we do not leverage the attention mechanism to train a neural model; rather explore the attention mechanism and analyze it for code.

Literature for analyzing the behavior of PLMs is limited. Karmakar et al. [21] studies the comprehension abilities of four different pre-trained models trained on code namely, BERT[10], CodeBERTa[48], CodeBERT[12] and GraphCodeBERT. They compare the performance of these models from the perspective of probing classifiers. Other studies [5, 9, 35, 49], show that identifiers are important code entities and can be used in the modeling of Transfoemr based models. This work presents the first study in software engineering, which analyses the multi-headed attention framework of BERT, which is not done previously [21]. Another difference between this work and [21] is the use of different probing classifiers. We use code clone detection and non-parametric probing for syntax tagging of code whereas in [21], it uses other probing tasks. Also, they compare the semantic and syntactic *performance of different PLMs*, but we explore the *semantic and syntax abilities of BERT with respect to what it learns about code constructs*.

Differences among our work and previous studies: Our study analyses where and how attention exists within the pre-trained models for code and what kind of information, both semantic and syntactic, is learned by the model. The current research recognizes the success and importance of pre-trained models; however, the software engineering research community has not focused on understanding the comprehension abilities of the PLMs for programming languages. The current studies do not generate any new embedding using identifiers, do not draw comparison with NLP, and do not study the attention mechanism of the models. A body of research in NLP focus on exploring the internal working of the PLMs on English. Our work is inspired and is closer to [8, 36, 39, 41]. The difference between our work and these studies is on different

domains: software engineering and NLP. We explore BERT’s semantic and syntactic behavior with respect to code constructs and assess our findings on code clone detection.

8 THREATS TO VALIDITY

Internal threats in this study can be related to inadequate training of the models. To eliminate this threat, we used the official implementation of BERT that is made available by its authors. The results obtained for training BERT on our dataset are similar to the original results of the BERT model reported by the authors of BERT, thus, ensuring the proper training of the model. Moreover, for CodeBERT, we use its official pre-trained weights provided by the authors, mitigating threats related to the training of CodeBERT.

External Threats relate to the dataset used in our work. We use a popular open-source dataset, CodeSearchNet [18]. The same dataset has been used in multiple PLM studies for code [12, 14, 28]. As the dataset is commonly used for pre-training PLMs, we anticipate low threats related to the choice of the dataset. Another threat is related to the chosen task. We only applied our analysis on code clone detection. Although more studies are needed, our results shed lights for other tasks as code clone detection uses both semantic and syntactic knowledge about the code. We only study Java, and the results might not be generalizable to other languages.

Conclusion Threats in our work refer to the threats caused by bias used in sampling code from the dataset. To mitigate this threat, we apply random sampling, and the experiments are repeated for both cased and uncased code sequences. The same results were obtained in both cases. Another bias could be related to the choice of model. To reduce this threat, we applied the identifier embedding for code clone detection on two different models, BERT and CodeBERT..

9 CONCLUSION AND FUTURE WORK

In this work, we studied the attention behavior of BERT for code and compared it to the attention behavior in NLP studies. We explored the relationship among the code constructs and found that identifiers have higher attention. An identifier attends towards itself more than other tokens; some of its attention value is made to its immediate left and right token rather than on the tokens farther from it. Based on this finding, we proposed to use the representation of identifiers in code clone detection, which requires the code semantic information. We showed this improves the clone detection results in all layers, and this finding is not restricted to BERT. The findings can be used to build smaller models which can then be used to employ the identifiers’ embeddings in the downstream task (e.g. by models with 6 layers instead of 12 layers), to attain the similar performance to that of larger models. A future direction could be evaluating the results on other software engineering tasks requiring code intelligence. It would also be interesting to confirm whether the combination of the identifiers’ and separators’ embeddings would generate meaningful embeddings or introduce noise for code.

ACKNOWLEDGMENTS

This research is support by a grant from Natural Sciences and Engineering Research Council of Canada RGPIN-2019-05175.

REFERENCES

- [1] Wasi Uddin Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. 2021. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- [2] Wasi Uddin Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. 2020. A Transformer-based Approach for Source Code Summarization. In *ACL*. 4998–5007. <https://www.aclweb.org/anthology/2020.acl-main.449/>
- [3] Uri Alon, Shaked Brody, Omer Levy, and Eran Yahav. 2019. code2seq: Generating Sequences from Structured Representations of Code. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019*.
- [4] Reem S. Alsuhaibani, Christian D. Newman, Michael J. Decker, Michael L. Collard, and Jonathan I. Maletic. 2021. On the Naming of Methods: A Survey of Professional Developers. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. 587–599. <https://doi.org/10.1109/ICSE43902.2021.00061>
- [5] Leonhard Applis, Annibale Panichella, and Arie van Deursen. 2021. Assessing Robustness of ML-Based Program Analysis Tools using Metamorphic Program Transformations. In *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 1377–1381.
- [6] I.D. Baxter, A. Yahin, L. Moura, M. Sant’Anna, and L. Bier. 1998. Clone detection using abstract syntax trees. In *Proceedings. International Conference on Software Maintenance (Cat. No. 98CB36272)*. 368–377. <https://doi.org/10.1109/ICSM.1998.738528>
- [7] Casey Casalnuovo, Kenji Sagae, and Premkumar T. Devanbu. 2019. Studying the Difference Between Natural and Programming Language Corpora. *Empirical Software Engineering* (2019).
- [8] Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. What Does BERT Look At? An Analysis of BERT’s Attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*.
- [9] Rhys Compton, Eibe Frank, Panos Patros, and Abigail Koay. 2020. Embedding Java Classes with code2vec. *Proceedings of the 17th International Conference on Mining Software Repositories* (Jun 2020). <https://doi.org/10.1145/3379597.3387445>
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [11] Dror Feitelson, Ayelet Mizrahi, Nofar Noy, Aviad Ben Shabat, Or Eliyahu, and Roy Sheffer. 2020. How Developers Choose Names. *IEEE Transactions on Software Engineering* (2020), 1–1. <https://doi.org/10.1109/TSE.2020.2976920>
- [12] Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, and Ming Zhou. 2020. CodeBERT: A Pre-Trained Model for Programming and Natural Languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, Online, 1536–1547. <https://doi.org/10.18653/v1/2020.findings-emnlp.139>
- [13] Shuzheng Gao, Cuiyun Gao, Yulan He, Jichuan Zeng, Lun Yiu Nie, and Xin Xia. 2021. Code Structure Guided Transformer for Source Code Summarization. *CoRR* abs/2104.09340 (2021). [arXiv:2104.09340](https://arxiv.org/abs/2104.09340) <https://arxiv.org/abs/2104.09340>
- [14] Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, Shujie Liu, Long Zhou, Nan Duan, Alexey Svyatkovskiy, Shengyu Fu, Michele Tufano, Shao Kun Deng, Colin Clement, Dawn Drain, Neel Sundaresan, Jian Yin, Daxin Jiang, and Ming Zhou. 2021. GraphCodeBERT: Pre-training Code Representations with Data Flow. In *International Conference on Learning Representations*.
- [15] Sakib Haque, Alexander LeClair, Lingfei Wu, and Collin McMillan. 2020. Improved Automatic Summarization of Subroutines via Attention to File Context. In *Proceedings of the 17th International Conference on Mining Software Repositories (Seoul, Republic of Korea) (MSR ’20)*. Association for Computing Machinery, New York, NY, USA, 300–310. <https://doi.org/10.1145/3379597.3387449>
- [16] Jacob A. Harer, Louis Y. Kim, Rebecca L. Russell, Onur Ozdemir, Leonard R. Kosta, Akshay Rangamani, Lei H. Hamilton, Gabriel I. Centeno, Jonathan R. Key, Paul M. Ellingwood, Erik Antelman, Alan Mackay, Marc W. McConley, Jeffrey M. Opper, Peter Chin, and Tomo Lazovich. 2018. Automated software vulnerability detection with machine learning. *arXiv e-prints*, Article arXiv:1803.04497 (feb 2018), arXiv:1803.04497 pages. [arXiv:1803.04497](https://arxiv.org/abs/1803.04497) [cs.SE]
- [17] Xing Hu, Ge Li, Xin Xia, David Lo, and Zhi Jin. 2018. Deep Code Comment Generation. In *Proceedings of the 26th Conference on Program Comprehension (Gothenburg, Sweden) (ICPC ’18)*. Association for Computing Machinery, New York, NY, USA, 200–210. <https://doi.org/10.1145/3196321.3196334>
- [18] Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. 2019. CodeSearchNet Challenge: Evaluating the State of Semantic Code Search. *arXiv e-prints*, Article arXiv:1909.09436 (Sept. 2019), arXiv:1909.09436 pages. [arXiv:1909.09436](https://arxiv.org/abs/1909.09436) [cs.LG]
- [19] Aditya Kanade, Petros Maniatis, Gogul Balakrishnan, and Kensen Shi. 2020. Learning and Evaluating Contextual Embedding of Source Code. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, Virtual, 5110–5121. <http://proceedings.mlr.press/v119/kanade20a.html>
- [20] Rafael-Michael Karampatsis and Charles Sutton. 2020. SCELMO: Source Code Embeddings from Language Models. *arXiv e-prints*, Article arXiv:2004.13214 (April 2020), arXiv:2004.13214 pages. [arXiv:2004.13214](https://arxiv.org/abs/2004.13214) [cs.SE]
- [21] Anjan Karmakar and Romain Robbes. 2021. What do pre-trained code models know about code? *CoRR* abs/2108.11308 (2021). [arXiv:2108.11308](https://arxiv.org/abs/2108.11308) <https://arxiv.org/abs/2108.11308>
- [22] J. Krinke. 2001. Identifying similar code with program dependence graphs. In *Proceedings Eighth Working Conference on Reverse Engineering*. 301–309. <https://doi.org/10.1109/WCRE.2001.957835>
- [23] Alexander LeClair, Siyuan Jiang, and Collin McMillan. 2019. A Neural Model for Generating Natural Language Summaries of Program Subroutines. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*.
- [24] Alexander LeClair, Siyuan Jiang, and Collin McMillan. 2019. A Neural Model for Generating Natural Language Summaries of Program Subroutines. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*.
- [25] Yi Li, Shaohua Wang, Tien N. Nguyen, and Son Van Nguyen. 2019. Improving Bug Detection via Context-Based Code Representation Learning and Attention-Based Neural Networks. *Proc. ACM Program. Lang.* 3, OOPSLA, Article 162 (oct 2019), 30 pages. <https://doi.org/10.1145/3360588>
- [26] Fang Liu, Ge Li, Yunfei Zhao, and Zhi Jin. 2020. Multi-task Learning based Pre-trained Language Model for Code Completion. In *2020 35th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 473–485.
- [27] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR* abs/1907.11692 (2019). [arXiv:1907.11692](https://arxiv.org/abs/1907.11692) <http://arxiv.org/abs/1907.11692>
- [28] Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin B. Clement, Dawn Drain, Daxin Jiang, Duyu Tang, Ge Li, Lidong Zhou, Linjun Shou, Long Zhou, Michele Tufano, Ming Gong, Ming Zhou, Nan Duan, Neel Sundaresan, Shao Kun Deng, Shengyu Fu, and Shujie Liu. 2021. CodeXGLUE: A Machine Learning Benchmark Dataset for Code Understanding and Generation. *CoRR* abs/2102.04664 (2021). [arXiv:2102.04664](https://arxiv.org/abs/2102.04664) <https://arxiv.org/abs/2102.04664>
- [29] Rabee S. Malik, Jibesh Patra, and Michael Pradel. 2019. NL2Type: Inferring JavaScript Function Types from Natural Language Information. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. 304–315. <https://doi.org/10.1109/ICSE.2019.00045>
- [30] Christopher D Manning, Kevin Clark, John Hewitt, Urvashi Khandelwal, and Omer Levy. 2020. Emergent linguistic structure in artificial neural networks trained by self-supervision. *Proceedings of the National Academy of Sciences* 117, 48 (2020), 30046–30054.
- [31] M.L. Menéndez, J.A. Pardo, L. Pardo, and M.C. Pardo. 1997. The Jensen-Shannon divergence. *Journal of the Franklin Institute* 334, 2 (1997), 307–318. [https://doi.org/10.1016/S0016-0032\(96\)00063-4](https://doi.org/10.1016/S0016-0032(96)00063-4)
- [32] Christian D. Newman, Reem S. AlSuhaibani, Michael J. Decker, Anthony Peruma, Dishant Kaushik, Mohamed Wiem Mkaouer, and Emily Hill. 2020. On the generation, structure, and semantics of grammar patterns in source code identifiers. *Journal of Systems and Software* 170 (2020), 110740. <https://doi.org/10.1016/j.jss.2020.110740>
- [33] Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H. Miller, and Sebastian Riedel. 2019. Language Models as Knowledge Bases?. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- [34] Michael Pradel and Koushik Sen. 2018. DeepBugs: A Learning Approach to Name-Based Bug Detection. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 147 (Oct. 2018), 25 pages. <https://doi.org/10.1145/3276517>
- [35] Md Rafiqul Islam Rabin, Nghi DQ Bui, Ke Wang, Yijun Yu, Lingxiao Jiang, and Mohammad Amin Alipour. 2021. On the generalizability of Neural Program Models with respect to semantic-preserving program transformations. *Information and Software Technology* 135 (2021), 106552.
- [36] Herbert Rubenstein and John B. Goodenough. 1965. Contextual Correlates of Synonymy. *Commun. ACM* 8, 10 (Oct. 1965), 627–633. <https://doi.org/10.1145/365628.365657>
- [37] Andrea Schankin, Annika Berger, Daniel V. Holt, Johannes C. Hofmeister, Till Riedel, and Michael Beigl. 2018. Descriptive Compound Identifier Names Improve Source Code Comprehension. In *2018 IEEE/ACM 26th International Conference on Program Comprehension (ICPC)*. 31–3109.
- [38] Jianhang Shuai, Ling Xu, Chao Liu, Meng Yan, Xin Xia, and Yan Lei. 2020. *Improving Code Search with Co-Attentive Representation Learning*. Association for Computing Machinery, New York, NY, USA, 196–207. <https://doi.org/10.1145/3387904.3389269>
- [39] Koustuv Sinha, Robin Jia, Dieuwke Hupkes, Joelle Pineau, Adina Williams, and Douwe Kiela. 2021. Masked Language Modeling and the Distributional Hypothesis: Order Word Matters Pre-training for Little. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.

- [40] Jeffrey Svajlenko, Judith F. Islam, Iman Keivanloo, Chanchal K. Roy, and Mohamad Mamun Mia. 2014. Towards a Big Data Curated Benchmark of Inter-project Code Clones. In *2014 IEEE International Conference on Software Maintenance and Evolution*. 476–480. <https://doi.org/10.1109/ICSME.2014.77>
- [41] Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. BERT rediscovers the classical NLP pipeline. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- [42] Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Sam Bowman, Dipanjan Das, and Ellie Pavlick. 2019. What do you learn from context? Probing for sentence structure in contextualized word representations. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=SJzSgnRcKX>
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *Advances in Neural Information Processing Systems*.
- [44] Jesse Vig. 2019. A Multiscale Visualization of Attention in the Transformer Model. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. Association for Computational Linguistics, Florence, Italy, 37–42. <https://doi.org/10.18653/v1/P19-3007>
- [45] Yao Wan, Zhou Zhao, Min Yang, Guandong Xu, Haochao Ying, Jian Wu, and Philip S. Yu. 2018. Improving Automatic Source Code Summarization via Deep Reinforcement Learning. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering (Montpellier, France) (ASE 2018)*. Association for Computing Machinery, New York, NY, USA, 397–407. <https://doi.org/10.1145/3238147.3238206>
- [46] Wenhua Wang, Yuqun Zhang, Zhengran Zeng, and Guandong Xu. 2020. TranS3: A Transformer-based Framework for Unifying Code Summarization and Code Search. *arXiv e-prints*, Article arXiv:2003.03238 (March 2020), arXiv:2003.03238 pages. arXiv:2003.03238 [cs.SE]
- [47] Martin White, Michele Tufano, Matías Martínez, Martin Monperrus, and Denys Poshyvanyk. 2019. Sorting and Transforming Program Repair Ingredients via Deep Learning Code Similarities. In *2019 IEEE 26th International Conference on Software Analysis, Evolution and Reengineering (SANER)*. 479–490. <https://doi.org/10.1109/SANER.2019.8668043>
- [48] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. HuggingFace’s Transformers: State-of-the-art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*.
- [49] Zhou Yang, Jieke Shi, Junda He, and David Lo. 2022. Natural Attack for Pre-trained Models of Code. In *2022 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*.
- [50] Jian Zhang, Xu Wang, Hongyu Zhang, Hailong Sun, and Xudong Liu. 2020. Retrieval-Based Neural Source Code Summarization. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering (Seoul, South Korea) (ICSE ’20)*. Association for Computing Machinery, New York, NY, USA, 1385–1397. <https://doi.org/10.1145/3377811.3380383>