LANGUAGE-AGNOSTIC REPRESENTATION LEARNING OF SOURCE CODE FROM STRUCTURE AND CONTEXT

Anonymous authors
Paper under double-blind review

ABSTRACT

Source code (*Context*) and its parsed abstract syntax tree (AST; *Structure*) are two complementary representations of the same computer program. Traditionally, designers of machine learning models have relied predominantly either on Structure or Context. We propose a new model, which jointly learns on Context and Structure of source code. In contrast to previous approaches, our model uses only language-agnostic features, i.e., source code and features that can be computed directly from the AST. Besides obtaining state-of-the-art on monolingual code summarization on all five programming languages considered in this work, we propose the first *multilingual* code summarization model. We show that jointly training on non-parallel data from multiple programming languages improves results on all individual languages, where the strongest gains are on low-resource languages. Remarkably, multilingual training only from Context does not lead to the same improvements, highlighting the benefits of combining Structure and Context for representation learning on code.

1 INTRODUCTION

Machine learning for code is an active and growing area of research which aims at building models that can learn semantically meaningful representations of programs. These embeddings can be used on downstream tasks, such as code generation, bug detection, or code summarization. We focus our work on two complementary data representations of programs: the source code (referred to as *Context* in this work), and the abstract syntax tree (AST; referred to as *Structure*). Traditionally, researchers and practitioners have decided to predominantly leverage either Structure or Context in their machine learning models. In this work, we show that jointly learning on Context and Structure improves representation learning on source code (see Fig. 1).

The source code representation naturally lends itself to models from natural language processing (NLP), e.g., long short-term memory networks (Hochreiter & Schmidhuber, 1997) (LSTM) or Transformers (Vaswani et al., 2017; Radford et al., 2019; Dai et al., 2019; Yang et al., 2019; Shaw et al., 2018). On the other hand, models leveraging the structure representations are typically based on graph neural networks (GNNs) (Kipf & Welling, 2017; Xu et al., 2019; Veličković et al., 2018; You et al., 2019; Hamilton et al., 2017; Li et al., 2015). While the AST representation makes the highly structured nature of source code explicit to the models, since most GNNs use the message-passing framework, their learned representations are inherently local and struggle to leverage long-range interactions.

Recently, Hellendoorn et al. (2020) have explored models that can leverage several representations, including both Structure and Context. Their Graph Relational Embedding Attention Transformer (GREAT) extends Shaw et al. (2018), which biases the self-attention computation in a localized way given the underlying graph structure. The language-specific representations used by GREAT include a combination of the data flow graph, control flow graph, syntactic edges (inspired by Allamanis et al. (2018)), etc. which require specialized pre-processing pipelines and static analysis tools in order to be obtained.

We propose the CODE TRANSFORMER, which combines distances computed on Structure and Context in the self-attention operation. In contrast to the localized treatment via edges described above, we make the full Structure accessible to the model at each layer by computing pairwise distances on

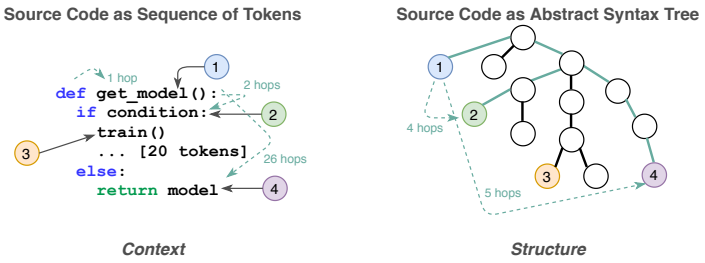

Figure 1: Context and Structure both encapsulate valuable information about source code. In this realistic example, token 1 and 4 are distant in the sequence of tokens (*Context*), but only 5 hops away when traversing the Abstract Syntax Tree (*Structure*). As such, a method that relies only on the sequence of tokens could neglect the relationship between a method name and its return variable. Conversely, token 1 and 2 showcase the opposite setting. Hence, unifying Structure and Context leads to a more powerful representation of source code.

the AST, such as shortest path lengths. To this end, we draw inspiration from the XLNet architecture (Yang et al., 2019), which uses relative distances instead of absolute positions in the attention computation. Importantly, all our features are *language-agnostic*, i.e., can easily be computed for any programming language based on the source code and AST.

We use two datasets comprising 5 different programming lanagues in total, and evaluate the representations learned by our model on the task of code summarization, where the model predicts a method's name based on its body. Besides setting the state-of-the-art on all five languages for single-language training, we also train the first *multilingual* model for code summarization. This is enabled by the fact that our model uses only language-agnostic features that can easily be obtained for any programming language. Remarkably, training our model on multiple programming languages substantially improves the performance on all languages. Moreover, multilingual training only from Context does not lead to the same improvements, highlighting the benefits of combining Structure and Context for representation learning on code.

2 Related Work

Machine Learning for Code. Early research learned language models on raw text data, e.g., (Wang et al., 2016; Raychev et al., 2014; Dam et al., 2016), providing evidence for the *naturalness assumption* (Hindle et al., 2012). For example, Allamanis et al. (2015) learned distributed representations of variables and methods, finding that they were indeed able to encode common semantic properties from the regularities present in source code. Alon et al. (2019b) also found evidence of semantic arithmetic in their embedding space, dubbed code2vec. These representations—and their variants like (Mou et al., 2016)—can then be used to predict sequences of identifier sub-tokens (Allamanis et al., 2015) or API calls (Acharya et al., 2007; Nguyen et al., 2017). They can be used as advanced auto-completion tools (Hindle et al., 2012; Bhoopchand et al., 2016), including for user-provided tokens like Variable Names (Raychev et al., 2014; Allamanis et al., 2014). These are useful for deobfuscating Android applications (Bichsel et al., 2016) for example.

Several works leverage structured graphical models for probabilistic models of source code, usually through parse trees (Maddison & Tarlow, 2014; Bielik et al., 2016). Unlike previous works where hand-crafted features were used as node features (Raychev et al., 2014) or as explicit semantic edges (Allamanis et al., 2018), our work does not augment the existing syntactic relationships between the different elements to enhance the predictive capabilities of the model. Other approaches (Alon et al., 2018; Li et al., 2017) also leverage the AST structure, but linearize the graph by first traversing it.

Learning representations of structured languages. While models of language have dramatically improved in their ability to learn structure (syntax) and semantic from scratch, it can be argued that directly providing the model with the underlying structure of the language can help with generalization (Battaglia et al., 2018), managing long-ranging dependencies (Tai et al., 2015), or representing the compositional aspect of natural language (Socher et al., 2013). Notably, tree structures have

shown promising results and inspired new architectures (Shen et al., 2019), including in the domain of source code (Fernandes et al., 2019), where the underlying syntax is directly available. Our work pursues this line of research, showing the benefits of explicitly integrating structural information as an inductive bias.

Graph Neural Networks. GNNs provide a powerful tool for machine learning on graphs, thanks to their ability to recursively incorporate information from neighboring nodes in the network (Battaglia et al., 2018), naturally capturing the graph structure simultaneously with the nodes' features. (Gori et al., 2005; Scarselli et al., 2008) are able to learn vector representations of nodes and graphs in an end-to-end fashion, encoding structural and feature information in the embedding space. Under this model, GNNs have achieved state-of-the-art performance across a variety of tasks, such as node classification (Kipf & Welling, 2017; Hamilton et al., 2017), link prediction (Zhang & Chen, 2018; Schlichtkrull et al., 2018), graph clustering (Defferrard et al., 2016; Ying et al., 2018) or graph classification (Ying et al., 2018; Dai et al., 2016; Duvenaud et al., 2015).

3 Integrating Structure and Context in the Code Transformer

Self-attention is the core operation powering the Transformer. It enables the model to selectively focus on relevant parts of the input. The matrix form equation for attention with a single head is

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

where $\boldsymbol{Q}, \boldsymbol{K} \in \mathbb{R}^{N \times d_k}$ and $\boldsymbol{V} \in \mathbb{R}^{N \times d_v}$. N is the number of input tokens, d_k the key dimension, and d_v the value dimension (typically we have $d_k = d_v$). The attention score of query $\boldsymbol{Q}_i$ and key $\boldsymbol{K}_j$ before softmax is

$$\boldsymbol{A}_{ij} = \boldsymbol{Q}_i^T \boldsymbol{K}_j = \boldsymbol{E}_i^T \boldsymbol{W}_q^T \boldsymbol{W}_k \boldsymbol{E}_j, \tag{2}$$

where $\boldsymbol{E}_i, \boldsymbol{E}_j \in \mathbb{R}^d$ are the d-dimensional embeddings of tokens i and j, and $\boldsymbol{W}_q, \boldsymbol{W}_k \in \mathbb{R}^{d_k \times d}$ are the query and key projection matrices, respectively.

Observe that Eq. (2) contains no assumption about potential structure in the input domain: in the attention operation we compute *all* dot products of query and key vectors equally, effectively viewing them as unordered sets of vectors. This means, however, that the model is oblivious to structured inputs (such as text or graphs) and therefore is unable to distinguish, for example, a variable name occurring as an argument and in the return statement of a method.

In NLP, it is common to bias Transformers towards sequential inputs by adding positional encodings to the token embeddings. These positional encodings are obtained by applying an encoding function $\phi : \mathbb{R} \to \mathbb{R}^d$ to each token's position p_i. These positional encodings make the information about the sequence of tokens available to the model. Eq. (2) becomes:

$$\boldsymbol{A}_{ij} = (\boldsymbol{E}_i + \phi(p_i))^T \boldsymbol{W}_q^T \boldsymbol{W}_k (\boldsymbol{E}_j + \phi(p_j)), , \tag{3}$$

which factorizes into

$$\boldsymbol{A}_{ij} = \underbrace{\boldsymbol{E}_i^T \boldsymbol{W}_q^T \boldsymbol{W}_k \boldsymbol{E}_j}_{\text{(a) } \boldsymbol{A}_{ij}^{\text{cc}}} + \underbrace{\boldsymbol{E}_i^T \boldsymbol{W}_q^T \boldsymbol{W}_k \phi(p_j)}_{\text{(b) } \boldsymbol{A}_{ij}^{\text{cp}}} + \underbrace{\phi(p_i)^T \boldsymbol{W}_q^T \boldsymbol{W}_k \boldsymbol{E}_j}_{\text{(c) } \boldsymbol{A}_{ij}^{\text{pc}}} + \underbrace{\phi(p_i)^T \boldsymbol{W}_q^T \boldsymbol{W}_k \phi(p_j)}_{\text{(d) } \boldsymbol{A}_{ij}^{\text{pp}}}. \tag{4}$$

We can interpret the terms (a)-(d) as follows. (a) $\boldsymbol{A}_{ij}^{\text{cc}}$ is the contribution from the 'match' between the content embeddings of tokens i and j; (b) $\boldsymbol{A}_{ij}^{\text{cp}}$ steers the attention towards certain positions based on the content of token i; (c) $\boldsymbol{A}_{ij}^{\text{pc}}$ biases towards content embeddings based on the position of token i; (d) $\boldsymbol{A}_{ij}^{\text{pp}}$ controls which positions should attend to which other positions.

In our model, we adopt the formulation of Dai et al. (2019); Yang et al. (2019). They modify Eq. (4) by replacing the *absolute* position encodings $\phi(p_i)$ with *relative* position encodings $\phi(r_{i \to j})$:

$$\boldsymbol{A}_{ij}^{\text{rel}} = \boldsymbol{E}_i^T \boldsymbol{W}_q^T \boldsymbol{W}_k \boldsymbol{E}_j + \boldsymbol{E}_i^T \boldsymbol{W}_q^T \boldsymbol{W}_r \phi(r_{i \to j}) + \boldsymbol{u}^T \boldsymbol{W}_k \boldsymbol{E}_j + \boldsymbol{v}^T \boldsymbol{W}_r \phi(r_{i \to j}), \tag{5}$$

where $r_{i \to j}$ is the relative distance from token i to token j in the sequence, $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^{d_k}$ are learnable bias vectors, and $\boldsymbol{W}_r$ is a key projection matrix for the relative distances. Besides fixing issues with absolute position encodings such as ambiguity when processing two sentences at a time, Eq. (5) enables native application of the powerful self-attention operation on domains such as graphs, where absolute coordinates are not available.

3.1 INTEGRATING SOURCE CODE AND AST REPRESENTATIONS OF PROGRAMS.

To enable the model to integrate information both the Context and Structure of programs, we modify Eq. (5) to be able to incorporate multiple different relations. To this end, we use one key projection matrix $\boldsymbol{W}_r^{(s)}$ *per relation* s, and sum their contributions in the raw attention score. This enables the CODE TRANSFORMER to combine information from multiple relations between tokens in the attention computation. Besides the token distance in the Context, we include pairwise relations based on the AST as described in the following. See Fig. 2 for a visualization of the Structure distances we use.

Shortest path length. We include the number of hops required to reach node j starting from node i and vice versa. Here, we treat the AST as an undirected graph, since otherwise most distances would be undefined: e.g., all other nodes in the AST would be unreachable from the leaves.

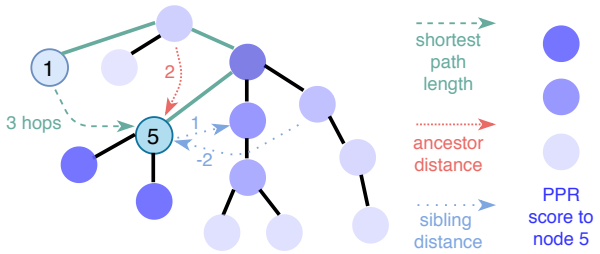

Figure 2: Structure distances used by our model.

Similar to the distance of two tokens on the source code sequence, the shortest-path length is a global distance. This makes the whole graph structure accessible to the model at each layer. In contrast, Hellendoorn et al. (2020) add bias terms to the attention computation only for edges (i.e. shortest-path distance of 1), which is a local operation that only exchanges information between immediate neighbors (similar to message passing in GNNs). The equivalent localized operation on the source code sequence would be to treat the sequence as a chain graph and only compute attention terms for neighboring tokens, which in turn highlights the benefit of non-local attention operations.

Ancestor distance. Since we treat the ASTs as undirected for the computation of the shortest-path length, we lose the direction information of the edges. To avoid this, we also include the distance on the ordered set of ancestors and descendants of a node in the AST (red arrow in Fig. 2). Again, we include number of (vertical) hops to avoid locality in the attention computation. For example, a node $r_{i \to j} = 2$ for "grand-children" j of i, and $r_{j \to i} = -2$ in the other direction.

Sibling distance. The neighbor sets in graphs are typically considered to be unordered, but in an AST, the order of children encodes their order of occurrence in the source code. To avoid information loss when encoding the AST, we further include the distance on the ordered set of siblings $\{v_i\}$ of a node, where we again avoid locality by encoding the number of hops, i.e. $r_{v_1 \to v_3} = 2$ and $r_{v_1 \to v_3} = -2$.

Personalized PageRank (Page et al., 1999) (PPR). PPR is a well-studied proximity measure has been shown to be very effective in learning with graphs (Klicpera et al., 2019a;b). PPR captures the local graph structure around a pair of nodes (i, j). E.g., if i has many neighbors, its PPR score for j will be low even when they are only few hops apart, which complements the purely hop-based distances described above.

3.2 EFFICIENT RELATIVE ATTENTION COMPUTATION.

Naïvely, we need to compute and materialize a tensor of dimension $N \times N \times d$ to hold all pairwise relative position encodings $\phi(r_{i \to j})$ in Eq. (5) , where N is the input length. This is prohibitive for fast GPU training. While for discrete distance values (e.g., sequence distance or shortest-path length on a graph) we only need to compute unique distance values occurring in the input, this does not generalize to continuous distances such as PPR. Therefore, we propose a constant-time approximation of the relational attention computation by grouping the values into $k \ll N^2$ bins. Since closer samples are typically more relevant for a query sample, we increase the bin widths exponentially with growing distance values. Throughout our experiments we have found the CODE TRANSFORMER to be relatively insensitive to the number of bins; we thus set $k = 32$ in our experiments.

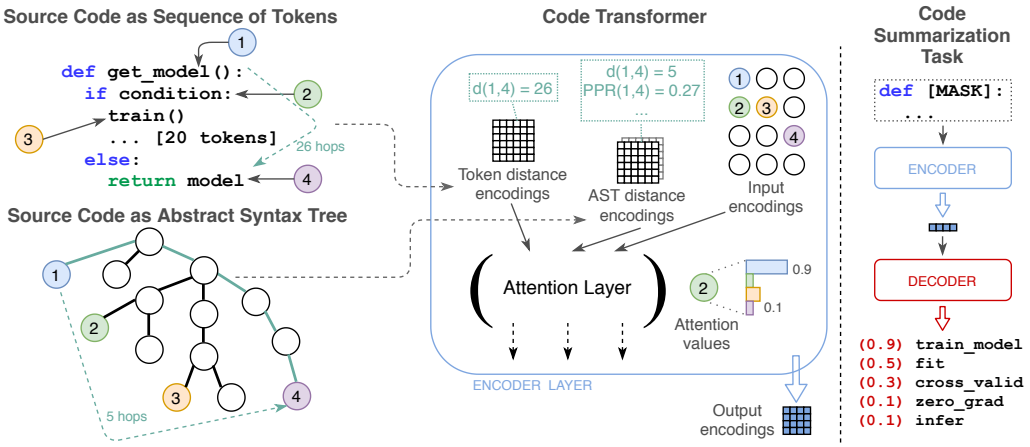

Figure 3: The CODE TRANSFORMER jointly leverages the sequence of tokens and the abstract syntax tree to learn expressive representations of source code. The output embeddings can be used for downstream tasks such as code summarization (right).

4 EXPERIMENTAL SETUP

Code summarization is one of the most popular tasks in machine learning for code. Given the body of a function, the task is to predict the function's name. As observed by Alon et al. (2019b) and Allamanis et al. (2016), this is a useful benchmark as method names in open-source projects tend to be precise and descriptive, and functions typically form complete logical units. See Fig. 3 (right) for a visual overview of the task. We use two complementary representations of programs: the source code as a sequence of tokens (Context) and the AST (Structure). As shown in Fig. 3 (left), tokens that are far away on the sequence may be very close on the AST and vice versa. In this task we make use of the CODE TRANSFORMER's ability jointly leverage both Structure and Context and show that it improves learning. Further, we show the benefit of using only language-agnostic features in our model by training the first multilingual model for code summarization.

Datasets. To highlight the benefit of only relying on language-agnostic representations such as source code and abstract syntax trees, we evaluate on challenging datasets in four programming languages introduced in the CodeSearchNet (CSN) Challenge (Husain et al., 2019): Python, Javascript, Go, and Ruby. Similar to Java-small, the datasets from CodeSearchNet have been carefully deduplicated by the creators to avoid data leakage from the training set, e.g., via copy-and-paste code.

We further evaluate on Java-small (Allamanis et al., 2016), a popular and challenging code summarization dataset. It contains 11 open-source Java projects. We use the split as in Alon et al. (2019a), where 9 of these projects are used for training, one for validation, and one for test. The dataset contains roughly 700K samples (function definitions). Moreover, we also experiment with pre-training our model on Java-medium and Java-large (Alon et al., 2019a) before fine-tuning on Java-small, making sure to avoid leakage by removing the test and validation projects of Java-small from the pre-training dataset. See Table 1 for a summary of the datasets we use in this work.

	Samples per partition		
Dataset	Train	Val.	Test
CSN-Python	412,178	23,107	22,176
CSN-Javascript	123,889	8,253	6,483
CSN-Ruby	48,791	2,209	2,279
CSN-Go	317,832	14,242	14,291
Java-small	691,974	23,844	57,088

Table 1: Dataset statistics.

Preprocessing. Each token of the source code is split into subtokens respective to code naming conventions, i.e., `get_TrainingData` is converted to `[get, training, data]`. Following Alon et al. (2019a) we use at most six subtokens for the method names, truncating longer function names if necessary. In addition to the tokenized source code we produce an AST for each method using the open-source AST parser Semantic[1]. We limit the vocabulary to subtokens with at least 100

[1] `https://github.com/github/semantic`

Model	Python			Javascript			Ruby			Go		
	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1
code2seq	35.79	24.85	29.34	30.18	19.88	23.97	23.23	10.31	14.28	52.30	43.43	47.45
GREAT	35.09	31.62	33.26	31.25	26.87	28.89	24.66	22.25	23.39	50.02	46.52	48.21
Ours w/o structure	**37.39**	32.01	34.49	33.33	26.86	29.75	29.88	**25.89**	**27.74**	51.79	47.58	49.60
Ours	36.41	**33.68**	**34.99**	**35.07**	**29.65**	**32.13**	**31.46**	24.50	27.55	**55.09**	**48.05**	**51.33**
Ours w/o structure (Mult.)	37.71	31.03	34.05	34.90	28.02	31.09	32.01	25.83	28.59	**56.26**	46.37	50.84
Ours (Multilanguage)	38.89	**33.84**	36.19	36.99	30.04	33.15	34.00	**29.01**	31.31	56.02	50.43	53.08
Ours (Mult. + Finetune)	39.85	32.82	35.99	**37.07**	29.85	33.07	**35.91**	27.79	**31.33**	55.63	51.11	**53.28**
Ours (Mult. + LM + Finetune)	**40.66**	33.66	**36.83**	36.27	**30.95**	**33.40**	32.62	28.38	30.35	54.18	**51.47**	52.79

Table 2: Code summarization results on the CSN dataset.

occurrences in the training set, and only consider snippets with 512 or fewer tokens (after removing punctuation). We refer the reader to the appendix for further details on the data preprocessing.

Pointer network. We add a pointer network (Vinyals et al., 2015) (as described in Fernandes et al. (2019)) to the decoders of all Transformer-based models. This enables them to enhance their predictions by pointing at positions in the input sequence. For instance, when predicting the method name `get_url`, the model can point directly to occurrences of the variable `url`. This often improves results for less frequent tokens, and even enables the model to predict tokens which are not in the vocabulary by pointing at their positions in the input.

Baselines. We compare with code2seq (Alon et al., 2019a), the Graph Relational Embedding Attention Transformer (GREAT) (Hellendoorn et al., 2020), and the BiLSTM+GNN→LSTM+Pointer model presented in Fernandes et al. (2019). Code2seq is a non-Transformer model and state of the art for code summarization using only AST information. GREAT is a recent Transformer model using the framework presented in (Shaw et al., 2018) to bias the attention via edges. In the original formulation, GREAT additionally uses hand-crafted, language-specific edges such as dataflow, 'computed from', or 'next lexical use' edges, which require specialized preprocessing and static analysis tools. While this approach of leveraging language-specific features can certainly improve results on specific tasks and programming languages, our goal is to have a flexible model that can be used on any programming language. Since the specialized preprocessing used by GREAT is proprietary and not public, we produce the results for GREAT using edges from the AST instead, i.e. it has access to the same information as our proposed model. Note that the preprocessing of Fernandes et al. (2019) is language specific, which is why we only compare with their results on Java-small.

5 RESULTS

5.1 MONOLINGUAL CODE SUMMARIZATION

CSN dataset. First, we study the performance (measured by F1 score) of our model and the baselines on the traditional setting, where training and evaluation are performed on a single programming language. The results are shown in the upper part of Table 2. The CODE TRANSFORMER (without multi-language training) substantially outperforms all other models on all but one language, highlighting the effectiveness of jointly learning from Structure and Context. The only exception is Ruby, where it performs on par with its Context-only variant. We attribute this to the fact that there are relatively few samples in the Ruby dataset, and that Ruby is an dynamically typed language, which could make the Structure less powerful for learning. Interestingly, the Context-only CODE TRANSFORMER outperforms GREAT on all languages. We attribute this to the fact that GREAT uses the Structure of the programs only in a localized way (see Sec. 3.1). Another noteworthy finding is that code2seq performs comparably to the Transformer-based baselines on Go. We hypothesize that ASTs are more informative on Go since it is a compiled and strongly typed language.

Java-small results. In Table 3 we present code summarization results on the Java-small dataset. Among all models equipped with a pointer network, the CODE TRANSFORMER (without pre-training) obtains state-of-the-art on code summarization, outperforming all baselines, including the previous state-of-the-art on Java-small proposed by Fernandes et al. (2019). Further, pre-training on Java-medium and Java-large substantially improves precision, recall, and F1 score after fine-tuning on Java-small. To avoid leakage, we removed the projects used in the validation and test splits of Java-small from the pre-training dataset.

Ablation study. We further perform ablations where we remove our model's access to the Context or Structure, also presented in Table 3.

With pointer network. We find that both ablations lead to a substantial drop in performance, highlighting the benefit of learning jointly from Structure and Context. Interestingly, the model without access to the Structure performs slightly better than the variant without Context. *Without pointer network.* We repeat the ablation on the variants without pointer network. Here, the variant without Context performs better than the variant without Structure, indicating that the pointer network helps to compensate for the lack of access to Structure. The Structure-only variant (w/o pointer net) of our model even outperforms the full variant in this scenario. Inspection of the results revealed that the Structure-only variant has better performance on longer method names, which have an outsize influence on the micro-F1 score used in this work.

Model	Prec.	Rec.	F1
Without pointer net			
code2seq	51.23	37.31	43.18
Ours w/o structure	50.84	45.66	48.11
Ours w/o context	51.85	47.08	49.35
Ours	50.51	46.96	48.67
With pointer net			
Fernandes et al. (2019)	-	-	51.4
GREAT	53.63	46.43	49.77
Ours w/o structure	55.50	46.09	50.36
Ours w/o context	54.54	45.46	49.59
Ours	54.58	49.43	51.83
Ours + Pretrain	**57.02**	**50.87**	**53.77**

Table 3: Results on Java-small and ablation study.

5.2 MULTILINGUAL CODE SUMMARIZATION

Setup. A key contribution of our proposed architecture is that it only uses language-agnostic features, i.e. the source code and features that can be directly computed from the AST. We use this fact to study the first *multi-language* code summarization model. To this end, we train the CODE TRANSFORMER jointly on Python, Javascript, Ruby, and Go. The shared sub-token vocabulary is the union of the individual language vocabularies, enabling us to evaluate the multi-language model on the individual languages and compare the results with the single-language models. As proposed by Conneau & Lample (2019), we add a learned language embedding to each input embedding.

Results. In the lower part of Table 2 we can see the results of training our CODE TRANSFORMER *jointly* on all four programming languages. Our multi-lingual variants substantially outperform the mono-lingual models on all languages. The strongest improvement is on Ruby, which is also the programming language with the smallest number of samples in the dataset. Fine-tuning on the individual languages after joint training on code summarization only has a marginal effect on performance, indicating that the multilingual objective is well-aligned with the individual languages. In the last row, we have a variant of our model where we pre-train on the multi-lingual masked language modeling task, followed by finetuning on code summarization on the individual languages.

Remarkably, the multi-language Context-only variant (i.e. without access to the Structure) performs substantially worse than the full multi-language variant. This highlights that Structure is crucial to exploit the commonalities of different programming languages.

Qualitative analysis of multilingual representations. Learning the CODE TRANSFORMER on multiple programming languages jointly provides us with embeddings in a shared representation space. In Fig. 4 we show a t-SNE (Maaten & Hinton, 2008) visualization of the ca. 40,000 snippets from the validation sets of four programming languages from the CSN dataset. We use the encoder output of the method name as a snippet's embedding. Note that the true method names are masked, i.e., inaccessible to the model.

While snippets from the same language tend to be grouped together, there are interesting intersections of the different programming languages. For example, we highlight all methods whose names start with the subtoken `parse` or `main`. We see that snippets starting with `parse` are predominantly in an intersection region of Python and Javascript. From these snippets, we display the cross-language pair with smallest Euclidean embedding distance in Fig. 5. Remarkably, both snippets are effectively the same method in Javascript and Python – it is worth reminding that the model has never seen any parallel data during training. On the other hand, snippets starting with `main` tend to lie at an intersectional region of Python, Javascript, and Go. In Table 5 in the appendix we show additional cross-lingual pairs with similar embeddings, including a failure case of a `main` function,

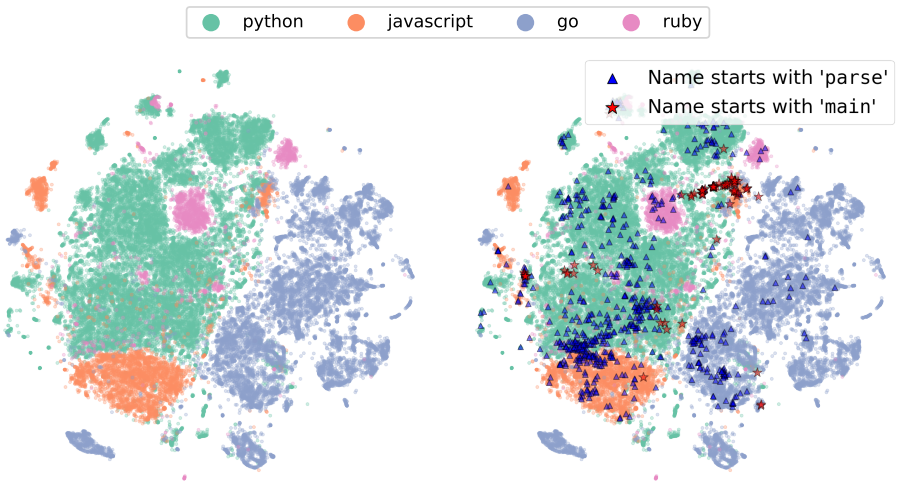

Figure 4: t-SNE visualization of the CODE TRANSFORMER's learned multilingual representations.

```
def parseBool(s):
  l = s.lower()
  if l in ("true", "t", "1"):
    return True
  if l in ("false", "f", "0"):
    return False
  raise Exception(
    "Unable to convert string '%s'"
    "to a boolean value" % s
  )
```

```
function jscoverage_getBooleanValue(s) {
  s = s.toLowerCase();
  if (s === 'false' || s === 'f'
    || s === 'no' || s === 'n'
    || s === 'off' || s === '0') {
    return false;
  }
  return true;
}
```

Figure 5: Example snippet starting with `parse` (left) and its best embedding match from other languages (right). Both methods parse an input string to convert it into a boolean value. Note that even though they are semantically very similar, their method names are not; nonetheless, their representations in the CODE TRANSFORMER reflect their semantic similarity.

where embedding distance is not representative of semantic similarity. We attribute this to the fact that we used the encoder output embedding of the *masked method name* – the representation used by the decoder to predict the method name – as a snippet's representation. Thus, snippets with completely different semantics (as is to be expected for very generic method names starting with `main`) have similar representations because they are predictive of the method name.

As another qualitative insight into the representations learned by the CODE TRANSFORMER we have found that the language embeddings of languages with similar roots in language design are close; see Table 4 in the appendix for the pairwise similarity matrix of the learned language embeddings.

6 CONCLUSION

We present the CODE TRANSFORMER, which learns jointly from Structure and Context of programs while only relying on language-agnostic features. Our model obtains state-of-the-art performance on code summarization on five different programming languages. In our ablation study, we find that learning jointly from Structure and Context improves learning. Besides these results for training on individual languages, the language-agnostic nature of our model allows us to train it *jointly* on multiple programming languages. The resulting multilingual model substantially outperforms its mono-lingual variant on all programming languages, setting the state of the art on each language. We observe the largest improvement from multilingual training on the language with fewest resources, indicating that multilingual training can improve learning for less widely used programming languages. Remarkably, multilingual training only from Context does not lead to the same improvements, highlighting the benefits of combining Structure and Context for representation learning on code.

REFERENCES

Mithun Acharya, Tao Xie, Jian Pei, and Jun Xu. Mining api patterns as partial orders from source code: From usage scenarios to specifications. In *Proceedings of the the 6th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering*, ESEC-FSE '07, pp. 25–34, New York, NY, USA, 2007. Association for Computing Machinery. ISBN 9781595938114. doi: 10.1145/1287624.1287630. URL https://doi.org/10.1145/1287624.1287630.

Miltiadis Allamanis, Earl T. Barr, Christian Bird, and Charles Sutton. Learning natural coding conventions. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*, FSE 2014, pp. 281–293, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450330565. doi: 10.1145/2635868.2635883. URL https://doi.org/10.1145/2635868.2635883.

Miltiadis Allamanis, Earl T. Barr, Christian Bird, and Charles Sutton. Suggesting accurate method and class names. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*, ESEC/FSE 2015, pp. 38–49, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450336758. doi: 10.1145/2786805.2786849. URL https://doi.org/10.1145/2786805.2786849.

Miltiadis Allamanis, Hao Peng, and Charles Sutton. A convolutional attention network for extreme summarization of source code. In *International conference on machine learning*, pp. 2091–2100, 2016.

Miltiadis Allamanis, Marc Brockschmidt, and Mahmoud Khademi. Learning to represent programs with graphs. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=BJOFETxR-.

Uri Alon, Meital Zilberstein, Omer Levy, and Eran Yahav. A general path-based representation for predicting program properties. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI 2018, pp. 404–419, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450356985. doi: 10.1145/3192366.3192412. URL https://doi.org/10.1145/3192366.3192412.

Uri Alon, Shaked Brody, Omer Levy, and Eran Yahav. code2seq: Generating sequences from structured representations of code. In *International Conference on Learning Representations*, 2019a. URL https://openreview.net/forum?id=H1gKYo09tX.

Uri Alon, Meital Zilberstein, Omer Levy, and Eran Yahav. code2vec: Learning distributed representations of code. *Proceedings of the ACM on Programming Languages*, 3(POPL):1–29, 2019b.

Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.

Avishkar Bhoopchand, Tim Rocktäschel, Earl Barr, and Sebastian Riedel. Learning python code suggestion with a sparse pointer network. *arXiv preprint arXiv:1611.08307*, 2016.

Benjamin Bichsel, Veselin Raychev, Petar Tsankov, and Martin Vechev. Statistical deobfuscation of android applications. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, pp. 343–355, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450341394. doi: 10.1145/2976749.2978422. URL https://doi.org/10.1145/2976749.2978422.

Pavol Bielik, Veselin Raychev, and Martin Vechev. Phog: probabilistic model for code. In *International Conference on Machine Learning*, pp. 2933–2942, 2016.

Alexis Conneau and Guillaume Lample. Cross-lingual language model pretraining. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 32*, pp. 7059–7069. Curran Associates, Inc., 2019. URL http://papers.nips.cc/paper/8928-cross-lingual-language-model-pretraining.pdf.

Hanjun Dai, Bo Dai, and Le Song. Discriminative embeddings of latent variable models for structured data. In *International conference on machine learning*, pp. 2702–2711, 2016.

Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019.

Hoa Khanh Dam, Truyen Tran, and Trang Pham. A deep language model for software code. In *FSE 2016: Proceedings of the Foundations Software Engineering International Symposium [The Conference]*, 2016.

Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in neural information processing systems*, pp. 3844–3852, 2016.

David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in neural information processing systems*, pp. 2224–2232, 2015.

Patrick Fernandes, Miltiadis Allamanis, and Marc Brockschmidt. Structured neural summarization. In *International Conference on Learning Representations*, 2019. URL `https://openreview.net/forum?id=H1ersoRqtm`.

Marco Gori, Gabriele Monfardini, and Franco Scarselli. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, pp. 729–734. IEEE, 2005.

William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *NIPS*, 2017.

Vincent J. Hellendoorn, Charles Sutton, Rishabh Singh, Petros Maniatis, and David Bieber. Global relational models of source code. In *International Conference on Learning Representations*, 2020. URL `https://openreview.net/forum?id=B1lnbRNtwr`.

Abram Hindle, Earl T. Barr, Zhendong Su, Mark Gabel, and Premkumar Devanbu. On the naturalness of software. In *Proceedings of the 34th International Conference on Software Engineering*, ICSE '12, pp. 837–847. IEEE Press, 2012. ISBN 9781467310673.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.

Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. Codesearchnet challenge: Evaluating the state of semantic code search. *arXiv preprint arXiv:1909.09436*, 2019.

Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.

Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann. Predict then propagate: Graph neural networks meet personalized pagerank. In *International Conference on Learning Representations (ICLR)*, 2019a.

Johannes Klicpera, Stefan Weißenberger, and Stephan Günnemann. Diffusion improves graph learning. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019b.

Jian Li, Yue Wang, Michael R Lyu, and Irwin King. Code completion with neural attention and pointer networks. *arXiv preprint arXiv:1711.09573*, 2017.

Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015.

Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.

Chris Maddison and Daniel Tarlow. Structured generative models of natural source code. In *International Conference on Machine Learning*, pp. 649–657, 2014.

Lili Mou, Ge Li, Lu Zhang, Tao Wang, and Zhi Jin. Convolutional neural networks over tree structures for programming language processing. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI'16, pp. 1287–1293. AAAI Press, 2016.

Trong Duc Nguyen, Anh Tuan Nguyen, Hung Dang Phan, and Tien N. Nguyen. Exploring api embedding for api usages and applications. In *Proceedings of the 39th International Conference on Software Engineering*, ICSE '17, pp. 438–449. IEEE Press, 2017. ISBN 9781538638682. doi: 10.1109/ICSE.2017.47. URL https://doi.org/10.1109/ICSE.2017.47.

Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9, 2019.

Veselin Raychev, Martin Vechev, and Eran Yahav. Code completion with statistical language models. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '14, pp. 419–428, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450327848. doi: 10.1145/2594291.2594321. URL https://doi.org/10.1145/2594291.2594321.

Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008.

Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In Aldo Gangemi, Roberto Navigli, Maria-Esther Vidal, Pascal Hitzler, Raphaël Troncy, Laura Hollink, Anna Tordai, and Mehwish Alam (eds.), *The Semantic Web*, pp. 593–607, Cham, 2018. Springer International Publishing. ISBN 978-3-319-93417-4.

Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. Self-attention with relative position representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pp. 464–468, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-2074. URL https://www.aclweb.org/anthology/N18-2074.

Yikang Shen, Shawn Tan, Alessandro Sordoni, and Aaron Courville. Ordered neurons: Integrating tree structures into recurrent neural networks. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=B1l6qiR5F7.

Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pp. 1631–1642, 2013.

Kai Sheng Tai, Richard Socher, and Christopher D Manning. Improved semantic representations from tree-structured long short-term memory networks. *arXiv preprint arXiv:1503.00075*, 2015.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph Attention Networks. *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=rJXMpikCZ.

Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. In *Advances in neural information processing systems*, pp. 2692–2700, 2015.

Song Wang, Devin Chollak, Dana Movshovitz-Attias, and Lin Tan. Bugram: Bug detection with n-gram language models. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering*, ASE 2016, pp. 708–719, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450338455. doi: 10.1145/2970276.2970341. URL https://doi.org/10.1145/2970276.2970341.

Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=ryGs6iA5Km.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pp. 5754–5764, 2019.

Zhitao Ying, Jiaxuan You, Christopher Morris, Xiang Ren, Will Hamilton, and Jure Leskovec. Hierarchical graph representation learning with differentiable pooling. In *Advances in neural information processing systems*, pp. 4800–4810, 2018.

Jiaxuan You, Rex Ying, and Jure Leskovec. Position-aware graph neural networks. *arXiv preprint arXiv:1906.04817*, 2019.

Muhan Zhang and Yixin Chen. Link prediction based on graph neural networks. In *Advances in Neural Information Processing Systems*, pp. 5165–5175, 2018.

	Python	Javascript	Go	Ruby
Python	1.00	0.43	0.42	0.79
Javascript	0.43	1.00	0.84	0.39
Go	0.43	0.84	1.00	0.38
Ruby	0.79	0.39	0.38	1.00

Table 4: Pairwise cosine similarities of the learned language embeddings of the CODE TRANSFORMER.

A APPENDIX

A.1 MULTILINGUAL REPRESENTATION ANALYSIS

In Table 4, we show the pairwise cosine similarities of the learned language embeddings of the CODE TRANSFORMER. We can see that the pairs Python-Ruby and Javascript-Go have similar language embeddings. This aligns well with roots of language design and common use cases of the languages.

Moreover, in Table 5, we show selected snippets starting with `is`, `main`, or `load` (left) and their best embedding matches from other languages (right).

A.2 DATA PREPROCESSING

A.2.1 TEXTUAL CODE SNIPPET PREPROCESSING

1. **Tokenize** code snippets with Pygments language-specific tokenizer
2. **Remove comments** (both multi-line, single-line and doc comments). The comment token types `pygments.token.Comment` and `pygments.token.Literal.String.Doc` that are generated by Pygments are used to identify comments
3. **Empty lines** are removed
4. **Hard coded strings and numbers** are replaced with a special `[MASK_STRING]` and `[MASK_NUMBER]` token
5. **Indentation** style of the code snippet is detected and whitespace characters at the beginning of a line are replaced with a single `[INDENT]` or `[DEDENT]` token when indentation changes
6. Tokens are further **split into sub tokens**, e.g., setBottomHeight $\rightarrow$ ['set', 'bottom', 'height']. Throughout our experiments, we use 5 input sub tokens. If a token consists of less than 5 sub tokens, the remaining spaces are filled with a special `[PAD]` token
7. Any remaining tokens that only consist of **white spaces are removed**. The only white space characters that are kept are line breaks '\n'
8. Snippets that are longer than `MAX_NUM_TOKENS` after punctuation tokens are removed are **discarded**. Throughout our experiments, we use `MAX_NUM_TOKENS = 512`. For the java-small dataset, 0.03% of the snippets are affected by this filtering step.
9. Any code snippets where the Pygments tokenizer **cannot parse a token** are also **discarded**.

A.2.2 STAGE 1 PREPROCESSING (GENERATION OF ASTS)

1. Stripped code snippets are used to **generate language-specific ASTs**. For Java, we use the AST parser from the **java-parser** project. The ASTs contain node types and source ranges.
2. Snippets that lead to an **AST parse error are discarded**.
3. We calculate a **mapping between tokens and nodes in the AST**. Every token is assigned to the node in the AST with tightest source range that still encompasses the source range of the token.

```javascript
function _isEqualArray(a, b) {
        if (a === b) {
            return true;
        }
        if ((a === undefined) ||
            (b === undefined)) {
            return false;
        }
        var i = a.length;
        if (i !== b.length){
            return false;
        }
        while (i--) {
            if (a[i] !== b[i]) {
                return false;
            }
        }
        return true;
    }
```

```go
func areSameFloat32Array(a, b []float32) bool {
  if len(a) != len(b) {
    return false
  }
  for i := 0; i < len(a); i++ {
    if a[i] != b[i] {
      return false
    }
  }
  return true
}
```

```javascript
function main() {
    var rawData = $('.HeaderTexture[data-login-
                    user-email]').data();
    if (rawData) {
        me = {
            name: rawData.loginUserName,
            mail: rawData.loginUserEmail
        };
        getCartId(function (cart) {
            me.cart = cart;
            injectMenu();
            refreshUsers(actions.updateUsers);
            listenOnChanges(onChange);
            listenOnOrderConfirm(onConfirm);
        });
    } else {
        callback('no user');
    }
}
```

```go
func TaskSayHello(t *tasking.T) {
  username := t.Flags.String("name")
  if username == "" {
    user, _ := user.Current()
    username = user.Name
  }

  if t.Flags.Bool("verbose") {
    t.Logf("Hello %s, the time now is %s\n",
           username, time.Now())
  } else {
    t.Logf("Hello %s\n", username)
  }
}
```

```perl
def _load_rule_file(self, filename):

        if not (os.path.exists(filename)):
            sys.stderr.write(
              "rflint: %s: No such file or "
              "directory\n" % filename
            )
            return
        try:
            basename = os.path.basename(filename)
            (name, ext) = os.path.splitext(basename)
            imp.load_source(name, filename)
        except Exception as e:
            sys.stderr.write(
              "rflint: %s: exception while "
              "loading: %s\n" % (filename, str(e))
            )
```

```go
func Backup(filename string) error {

  info, err := os.Stat(filename)
  if err != nil {
    if os.IsNotExist(err) {
      return nil
    }
    return err
  }
  if info.Size() == 0 {
    return nil
  }

  files, err := filepath.Glob(
    filename + _BACKUP_SUFFIX
  )
  if err != nil {
    return err
  }

  numBackup := byte(1)

  if len(files) != 0 {
    lastFile := files[len(files)-1]
    numBackup = lastFile[len(lastFile)-2] + 1
    if numBackup > '9' {
      numBackup = '1'
    }
  } else {
    numBackup = '1'
  }

  return Copy(filename,
              fmt.Sprintf("%s+%s~", filename,
                          string(numBackup)))
}
```

Table 5: Selected snippets starting with `is`, `main`, or `load` (left) and their best embedding matches from other languages (right).

To find such a node, we originally intended to make use of the assumption that source ranges of child nodes do not overlap. Then, one could easily find the node with smallest encompassing source range by greedily selecting at every layer in the AST the child that encompasses the token's source range (there can only be at most one child that fulfills this). However, this assumption does not hold for all ASTs. As an approximation, we greedily select the child node with the shorter source range in case there were multiple child nodes with encompassing source ranges. This approximation seems to be sufficient in our case, and limits runtime as we do not have to consider multiple paths in the AST. It is also sufficient to stop when no child node encompasses the source range of the token, as in ASTs the souce ranges of child nodes are always contained in the source ranges of their parent

A.2.3 STAGE 2 PREPROCESSING (CALCULATION OF DISTANCE MATRICES)

1. Tokens are **vocabularized**. Any token occurring less than 100 times in the training set is **replaced by an `<unk>`** token
2. We calculate multiple **pair-wise relations** between nodes in the AST:
   - Personalized Page Rank (PPR)
     We interpret the negative logarithm of PPR as a distance. We use a teleport probability of $\alpha = 0.15$ and a threshold of e^{-5}, i.e., anything with $-\log PPR > 5$ is considered unreachable
   - Shortest path length between two nodes
   - Ancestor shortest paths (bidirectional). That is, the parent has an ancestor shortest path distance of 1 to all its children and the child hase a distance of -1 to its parents.
   - Next sibling shortest paths (bidirectional, analogous to the ancestor shortest paths)

   Note that the ancestor shortest paths and next sibling shortest paths are required because treating the AST as a normal graph leads to ambiguity. In a graph, the neighbors of a node have no ordering; however in the AST, the order of the children of a node reflects their order in the code. Therefore, we explicitly include the next sibling shortest paths. The ancestor shortest paths would not be required if we treated the AST as a directed graph; in this case, however, a leaf node could not reach any other node in the AST, and therefore both PPR and shortest path length are not useful in this case. Therefore, we model the AST as undirected and inject the ancestor / child edges to avoid ambiguity.
3. **Distance values are binned** into 32 bins using area-based exponential binning with a growth factor of 1.3, i.e., the area of a bin's rectangle (x: bin range, y: number of values in bin) will be approximately 1.3 times bigger for the next bin (going away from the bin that contains the zero value). Additionally, for discrete distance measures (such as sequence distance or shortest path length), we hard-code 9 values around 0 to have their own bins. For instance on the sequence distance the values $-4, -3, \ldots, 4$ have their individual bins, and around those values we employ the exponential binning.

A.3 INPUT EMBEDDINGS TO THE MODEL

Besides its five subtokens (e.g., `['get', 'data', '[PAD]', '[PAD]', '[PAD]']`), each input token has a token type (coming from the Pygments tokenizer) and an AST node type. The AST node type is the type of the node assigned to each respective token, as described in Section A.2.2. We concatenate the embeddings of the five subtokens, the token type, and the AST node type. Then, we apply a linear layer (without activation function) to project down to the model's embedding dimension.

A.4 EXPERIMENTAL SETUP

Table 6 shows hyperparameters of our models for code summarization. For all our experiments, we use a Transformer Decoder with one layer and teacher forcing to generate 6 output sub tokens. We also employ label smoothing of 0.1. As optimizer, we use Adam with a learning rate of $8e^{-5}$ and weight decay of $3e^{-5}$. Batch size during training is 8 with a simulated batch size of 128 achieved

(a) CODE TRANSFORMER

Hyperparameter	Value
Activation	GELU
Input Nonlinearity	tanh
Num. layers	3
d	1024
d_{FF}	2048
p_{dropout}	0.2
Num. heads	8

(b) GREAT (Hellendoorn et al., 2020)

Hyperparameter	Value
Activation	GELU
Num. layers	3
d	1024
d_{FF}	2048
p_{dropout}	0.2
Num. heads	8

Table 6: Code Summarization hyperparameters

```
public Summation next() {
    return parts[i++];
}
```

Model	Prediction
GREAT	get x map
code2seq	get parts
Ours w/o structure	get
CODE TRANSFORMER	get next
Ground Truth	next

Table 8: The CODE TRANSFORMER is the only model to correctly identify the notion of getting the *next* entry.

by gradient accumulation.

Apart from comparing the CODE TRANSFORMER to baselines, we performed the following hyperparameter comparisons and ablation studies:

- CODE TRANSFORMER (structure-only)
  Using only AST information as input, i.e., masking all tokens that do not correspond to a leaf of the AST, and removing the token distance as a relation to be used by the model. Further, token types are not fed into the model.
- CODE TRANSFORMER (context-only)
  Here, we do not include any information on the AST (i.e. node types and distances on the AST). This is effectively the XLNet backbone plus encoding of the token type returned by the tokenizer.
- CODE TRANSFORMER (Max-Dist.)
  Applying a Max Distance Mask of 5 to the shortest paths distance (i.e., model cannot see a node that is more than 5 hops away no matter how small the other distances are). Early results showed that, as expected, results deteriorate substantially when limiting our model's receptive field. Hence, we do not include these results in this work.
- Using 16 and 64 bins instead of 32 bins. This had no noticeable effect on performance.

A.5 CODE SUMMARIZATION EXAMPLES

In the Tables 8, 9, 10, 11, 12, 13, 14 and 15 we present example functions from the Java-small dataset along with the different models' predictions for the function name.

```
private Path findCacheFile(Path[] cacheFiles, String fileName) {
    if (cacheFiles != null && cacheFiles.length > 0) {
        for (Path file : cacheFiles) {
            if (file.getName().equals(fileName)) {
                return file;
            }
        }
    }
    return null;
}
```

Model	Prediction
GREAT	get path
code2seq	find file
Ours w/o structure	get file
CODE TRANSFORMER	find cache
Ground Truth	find cache file

Table 9: The CODE TRANSFORMER is the only model to both recognize that the task is to *find* a file as well as the fact that it is about the cache. However, it did not correctly predict the *file* part of the method name.

```
public int compare(Pair<LoggedJob, JobTraceReader> p1,
                   Pair<LoggedJob, JobTraceReader> p2) {
    LoggedJob j1 = p1.first();
    LoggedJob j2 = p2.first();
    return (j1.getSubmitTime() < j2.getSubmitTime()) ? -1
            : (j1.getSubmitTime() == j2.getSubmitTime()) ? 0 : 1;
    }
```

Model	Prediction
GREAT	run
code2seq	get submit time
Ours w/o structure	compare
CODE TRANSFORMER	compare
Ground Truth	compare

Table 10: The CODE TRANSFORMER and the its context-only variant are the only models correctly recognizing the 'compare' template in the method body.

```
public static MNTPROC fromValue(int value) {
    if (value < 0 || value >= values().length) {
        return null;
    }
    return values()[value];
}
```

Model	Prediction
GREAT	get value
code2seq	get value
Ours w/o structure	to
CODE TRANSFORMER	from value
Ground Truth	from value

Table 11: The CODE TRANSFORMER is the only model to recognize that the snippet is similar to a static factory method which is often preceded with *from*.

```
private Iterable<ListBlobItem> listRootBlobs(String aPrefix,
                                             boolean useFlatBlobListing,
                                             EnumSet<BlobListingDetails> listingDetails,
                                             BlobRequestOptions options,
                                             OperationContext opContext)
                                        throws StorageException, URISyntaxException {
    CloudBlobDirectoryWrapper directory = this.container.getDirectoryReference(aPrefix);
    return directory.listBlobs(null, useFlatBlobListing,
                               listingDetails, options, opContext);
}
```

Model	Prediction
GREAT	list blobs
code2seq	list blobs
Ours w/o structure	list blobs
CODE TRANSFORMER	list blobs by prefix
Ground Truth	list root blobs

Table 12: All models could correctly identify the *listBlobs()* call in the return statement. However, the CODE TRANSFORMER additionally comprehended that the specified prefix is quite important.

```
private static void dumpOpCounts(EnumMap<FSEditLogOpCodes, Holder<Integer>> opCounts) {
    StringBuilder sb = new StringBuilder();
    sb.append("Summary of operations loaded from edit log:\n  ");
    Joiner.on("\n  ").withKeyValueSeparator("=").appendTo(sb, opCounts);
    FSImage.LOG.debug(sb.toString());
}
```

Model	Prediction
GREAT	append
code2seq	add
Ours w/o structure	log
CODE TRANSFORMER	log op counts
Ground Truth	dump op counts

Table 13: Only the CODE TRANSFORMER could correctly identify that it is the *op counts* that should be logged.

```
static String execCommand(File f, String... cmd) throws IOException {
    String[] args = new String[cmd.length + 1];
    System.arraycopy(cmd, 0, args, 0, cmd.length);
    args[cmd.length] = f.getCanonicalPath();
    String output = Shell.execCommand(args);
    return output;
}
```

Model	Prediction
GREAT	get canonical path
code2seq	exec
Ours w/o structure	get output
CODE TRANSFORMER	exec command
Ground Truth	exec command

Table 14: Only the CODE TRANSFORMER and code2seq could identify that the relevant part of the method is concerned with executing a command instead of returning something.

```
protected void subView(Class<? extends SubView> cls) {
    indent(of(ENDTAG));
    sb.setLength(0);
    out.print(sb.append('[').append(cls.getName()).append(']').toString());
    out.println();
}
```

Model	Prediction
GREAT	print
code2seq	print
Ours w/o structure	print
CODE TRANSFORMER	print sub view
Ground Truth	sub view

Table 15: Only the CODE TRANSFORMER was able to link the *print* functionality to the object that should be printed, which can only be inferred from the object's class in the method parameters.