

Improving Code Summarization with Block-wise Abstract Syntax Tree Splitting

Chen Lin, Zhichao Ouyang, Junqing Zhuang, Jianqiang Chen, Hui Li*, Rongxin Wu
Xiamen University

Xiamen, China

{chenlin, hui, wurongxin}@xmu.edu.cn, {ouyangzhichao, zhuangjq, jqchen}@stu.xmu.edu.cn

Abstract—Automatic code summarization frees software developers from the heavy burden of manual commenting and benefits software development and maintenance. Abstract Syntax Tree (AST), which depicts the source code’s syntactic structure, has been incorporated to guide the generation of code summaries. However, existing AST based methods suffer from the difficulty of training and generate inadequate code summaries. In this paper, we present the Block-wise Abstract Syntax Tree Splitting method (BASTS for short), which fully utilizes the rich tree-form syntax structure in ASTs, for improving code summarization. BASTS splits the code of a method based on the blocks in the dominator tree of the Control Flow Graph, and generates a split AST for each code split. Each split AST is then modeled by a Tree-LSTM using a pre-training strategy to capture local non-linear syntax encoding. The learned syntax encoding is combined with code encoding, and fed into Transformer to generate high-quality code summaries. Comprehensive experiments on benchmarks have demonstrated that BASTS significantly outperforms state-of-the-art approaches in terms of various evaluation metrics. To facilitate reproducibility, our implementation is available at <https://github.com/XMUDM/BASTS>.

Index Terms—Code Summarization, Code Splitting, Abstract Syntax Tree

I. INTRODUCTION

Human-readable code summaries (i.e., code comments) are the key to enhance the comprehension of source code in software development and maintenance [1], [2]. As software projects grow rapidly in scale and complexity, manual commenting becomes costly for developers. There has been increasing interest in *automatic code summarization* [3], which automatically generates short and accurate textual descriptions for software entities and frees software developers from the tremendous workload of writing comprehensible comments.

Early works on code summarization select a subset of the statements and keywords from the code to generate comments [3]. However, the grammar of programming languages strictly follows specific syntaxes, which is not easy to model using code tokens only. It is critical to encode both the meaning of each code token and the syntax structure when generating code summarization. Abstract Syntax Tree (AST), which depicts the hierarchical abstract syntactic structure of source code, is therefore commonly harnessed together with code tokens to improve the quality of code summarizations [4]–[6]. Nevertheless, it is infeasible to directly adopt the methods studying the syntax tree in Natural Language

Processing (NLP) and apply tree-aware neural networks (e.g., Tree-LSTM [7]) in the code summarization task. The reason is that ASTs are quite different compared to NLP syntax trees. The size of an AST in terms of tree depth is very large, leading to a heavy encoder which is difficult and slow to train. For instance, the maximum AST depth of the two datasets in our experiments (Sec. VI) is 76 (Java) and 112 (Python), which is rare for NLP syntax trees and weakens the capability of neural networks to capture the complex semantics [8]. Graph Neural Network (GNN) [9] based methods have the similar problem of the heavy encoder and their results do not show significant improvements over other methods in code summarization [10].

To overcome the challenge of training over the entire AST, a great number of methods are proposed. But they all produce inadequate summaries. We now use the example from the Java dataset in our experiments to illustrate their limitations. In Fig. 1, reference (i.e., gold standard summary) fully reveals the hierarchical syntax in the AST and the long-term dependence in the source code. For example, in the hierarchical structure of the AST, “*timemillis*” is linked to “*expire*” to form a sub-clause, where the sub-clause is in turn linked to a connection status in an “*if*” clause. The “*if*” clause triggers the “*close*” action via the long-term dependence. Accordingly, in the reference, “*timemillis*” in the “*idle*” clause complements “*connection*”. It also complements “*close*” with “*connection (if it has) idled for timemillis*”. Next, let us see the results from contemporary code summarization methods which can be roughly divided into three categories:

- 1) Most works serialize ASTs and then adopt a sequential neural model. Specifically, an AST can be serialized as a single sequence of nodes [4], [11]–[13] via traversal, or a collection of paths [5], [6] (i.e., several sequences of nodes). Serialized ASTs can be easily modeled, as modeling sequential natural language tokens [14], by sequential neural models such as Recurrent Neural Network (RNN) [15], Long Short Term Memory (LSTM) [16], Gated Recurrent Unit (GRU) [17] or Transformer [18]. However, such a strategy encodes only the linear order syntax structure and inevitably loses *hierarchical syntax information*. In Fig. 1, AST-attndgru [11], Hybrid-DeepCom [12] and Rencos [13] which learn from serialized ASTs, are unable to capture the relations (including “*connection is expired*” or “*close the connection*”) correctly and perform even

* Hui Li is the corresponding author.

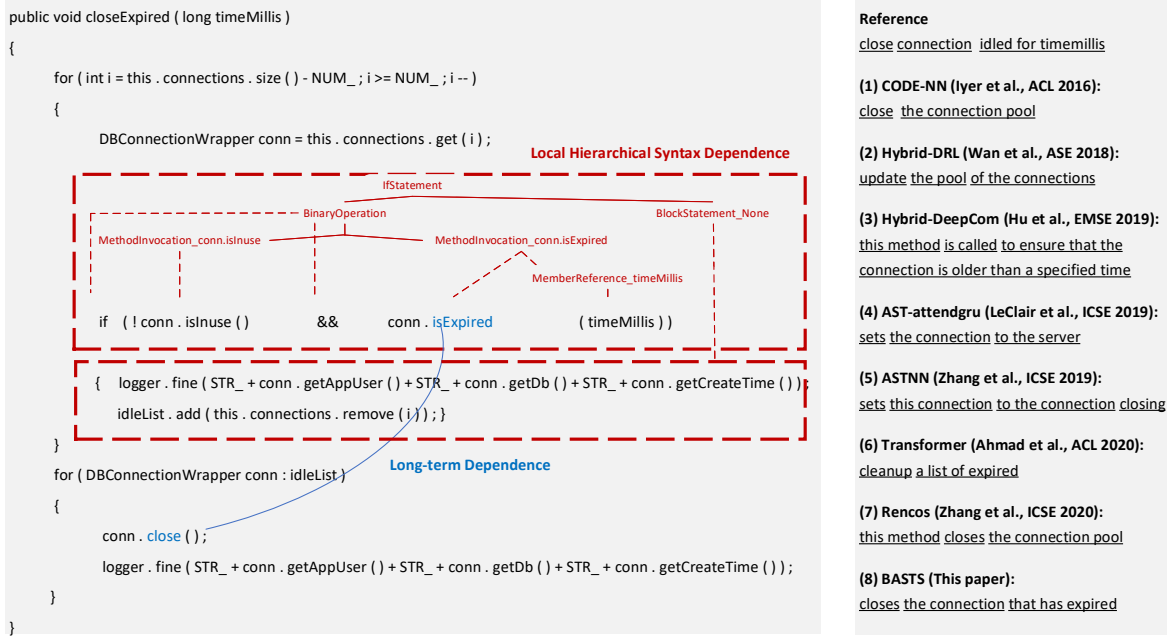


Fig. 1: Outputs of different automatic code summarization methods for test sample #1445 in the Java dataset

- worse than Code-NN [19] which does not consider ASTs.
- 2) A few works simplify the tree structure of ASTs, e.g., transform an AST to a binary tree [20]. However, they change the original semantics of code and bring information loss [21]. In Fig. 1, Hybrid-DRL [20], which learns from binarized AST, is unable to capture the long-term dependence between “connection is expired” and “close the connection” as well as the relation between “close” and “connection”.
 - 3) Last direction is most related to our work. ASTNN [21] splits an AST into small statement trees, and each tree consists of AST nodes of one statement. The statement tree is further encoded by Recursive Neural Network (RvNN) [22]. This way, the difficulty of training is reduced. Nevertheless, such a strategy forces RvNN to focus on one statement each time, as it only see one statement in each split AST. The long-term dependence, which is beyond the boundary of one statement, cannot be captured by RvNN only. If the subsequent sequential neural model fails to fuse multiple statement trees well, the generated summary will solely contain information from one statement. For instance, in Fig. 1, ASTNN fails to capture the dependence between “connection is expired” and “close the connection”. More analysis can be found in Sec. VI-C.

In this paper, we propose the **Block-wise Abstract Syntax Tree Splitting (BASTS)** to address aforementioned issues. Specifically, we split the code of a method based on the *blocks* in the dominator tree [23] of the Control Flow Graph of the method [24], instead of statements in the code [13], to capture the information beyond the boundary of one statement. After splitting, for each split code, we generate a *split AST*, which is modeled by a Tree-LSTM. Furthermore, inspired by fine-

tuning a pre-trained model on specific tasks in NLP [25], we adopt pre-training to learn split AST representations by predicting next split AST in the dominator tree. A high-quality comment requires understanding of both local dependence and long-term dependence. Therefore, we combine the split AST representations with the source code representations in the Transformer architecture [18] which has the ability to model both types of dependences. As shown in Fig. 1, only BASTS accurately recovers all relations in the reference.

The proposed BASTS framework has three advantages:

- 1) Through code splitting, each AST is reduced to an affordable tree propagation space for building the Tree-LSTM, and each block still retains sufficient information. Thus, BASTS is efficient and scalable for large programs.
- 2) Pre-training involves only the syntax representations. Hence, by first pre-training and then fine-tuning the syntax representations with token representations, BASTS saves computational resources, compared to building both representations from scratch. Moreover, the pre-training is unsupervised, which is endowed with the potential to handle sparse training data without sufficient code comments.
- 3) BASTS is effective since it learns both long-term dependencies on tokens (i.e., with the attention mechanism of Transformer) and local hierarchical dependencies on syntax structures (i.e., with Tree-LSTM in the local context).

In summary, the contributions of this paper are:

- 1) We present BASTS which incorporates normal tree structure of AST rather than serialized AST or binary AST to avoid information loss in code representation learning.
- 2) We improve automatic code summarization by using block-wise AST splitting, pre-training and Transformer to effectively learn local and hierarchical syntax dependences.

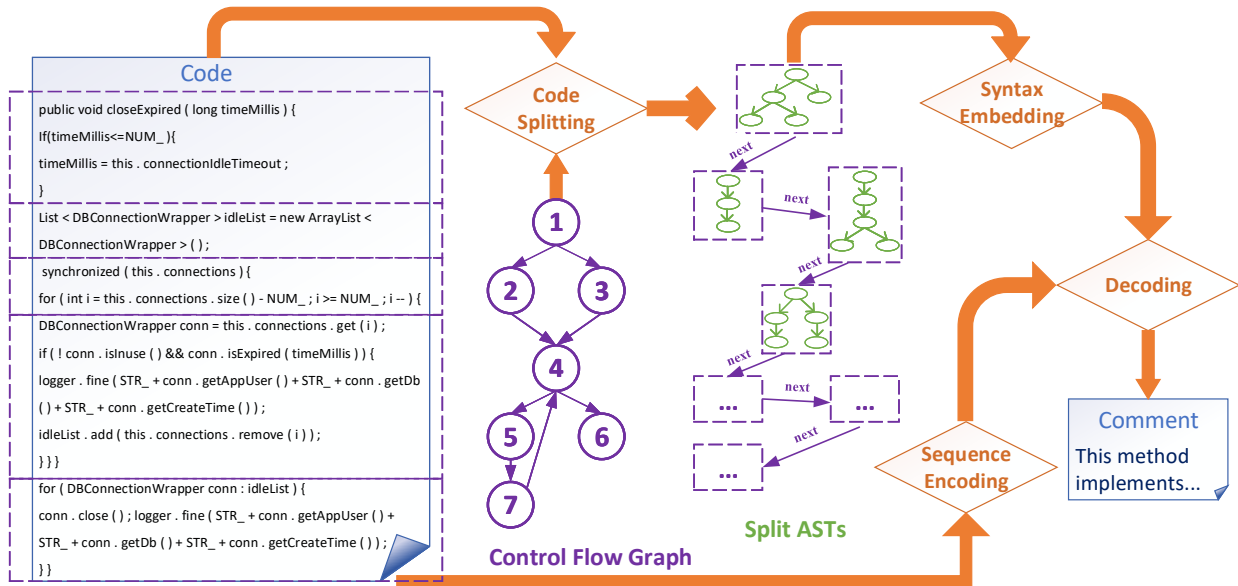


Fig. 2: Overview of BASTS. The CFG and split ASTs are simplified for illustration purpose.

3) We conduct comprehensive experiments on benchmarks for code summarization to verify the effectiveness of BASTS.

The rest of the paper is organized as follows. We present an overview of BASTS in Sec. II. We illustrate the block-wise code splitting in Sec. III. We describe in detail the neural models of pre-training syntax representations in Sec. IV and code summarization in Sec. V. We report and analyze the experimental results in Sec. VI. We elaborate on related work in Sec. VII. Finally, we conclude our work in Sec. VIII.

II. OVERVIEW

Given a program (e.g., code fragment for a method or a class), the aim of code summarization is to generate a short paragraph of concise and coherent comments to describe the function of the program. Hereafter, unless stated otherwise, we use lower-case letters for variables, upper-case letters for constants, lower-case bold letters for vectors, upper-case bold-face letters for matrices, upper-case calligraphic letters for sets, superscripts for sample index, and subscripts for vector index. To be specific, we are given a set of source programs and comments $\mathcal{N} = \{\mathbf{x}^n, \mathbf{y}^n\}, 1 \leq n \leq N$, where N is the number of training instances. Each source program is a sequence of tokens, i.e. $\mathbf{x}^n = \langle \mathbf{x}_1^n, \dots, \mathbf{x}_{D_n}^n \rangle$, where D_n is the number of tokens for the n -th sample. In the training set, each source program is paired with a comment, which is also a sequence of words, i.e., $\mathbf{y}^n = \langle \mathbf{y}_1^n, \dots, \mathbf{y}_{S_n}^n \rangle$, where S_n is the number of words for the n -th sample. The tokens are from the programming language vocabulary and the words are from the natural language vocabulary, i.e., $\mathbf{x}_d^n \in \mathcal{O}, 1 \leq d \leq D_n$ and $\mathbf{y}_s^n \in \mathcal{W}, 1 \leq s \leq S_n$.

BASTS consists of three components. As shown in Fig. 2, we first pre-process the program by code splitting which is performed according to the Control Flow Graph (CFG), and we generate a set of split ASTs on the split code. Then, in *tree encoding*, we pre-train the representations for each

split AST to obtain *syntax embedding*, i.e., hierarchical syntax representations of the split AST. Finally, the summarization model is trained based on *decoding* input, including the pre-trained syntax embedding in tree encoding, and the *sequence encoding* of source code \mathbf{x}_n , to output comment \mathbf{y}^n .

III. CODE SPLITTING

The goal of code splitting is to facilitate encoding local hierarchical syntax structure. We expect that each split code after the splitting would provide concise and local context to locate tree-form syntax dependence among code tokens. In this work, we choose to split the code of a method based on the blocks in the dominator tree [23] of the CFG of the method [24]. Following steps show how to generate split code:

- 1) Construct CFG to capture control flow relationships among statements. In the following, we will omit the sample index and use $G = (\mathcal{U}(G), \mathcal{E}(G))$ to denote the CFG for the current training sample. A CFG is a directed graph where each node $u \in \mathcal{U}(G)$ represents a statement, and each edge $e \in \mathcal{E}(G) = (u \rightarrow u')$ represents the control flow from statement u to u' . Note that the start node and end node of the CFG are both virtual which do not relate to any real statements.
- 2) Construct the dominator tree on the CFG. A dominator tree is a directed graph $DT = (\mathcal{U}(DT), \mathcal{E}(DT))$, where $\mathcal{U}(DT) = \mathcal{U}(G)$, and $\mathcal{E}(DT) \subset \mathcal{E}(G)$. Each edge is connected, i.e., $e \in \mathcal{E}(DT) = (u \rightarrow u')$ if and only if u dominates u' . u dominates u' if and only if every path in CFG from the start node to u' goes through u .
- 3) Partition the dominator tree into the blocks as split code. Each block in the dominator tree is a set of consecutive nodes containing no branches except at the very end. Thus we remove every edge $e \in \mathcal{E}(DT) = (u \rightarrow u')$ if u' has more than one in-coming edges, or u has more than one

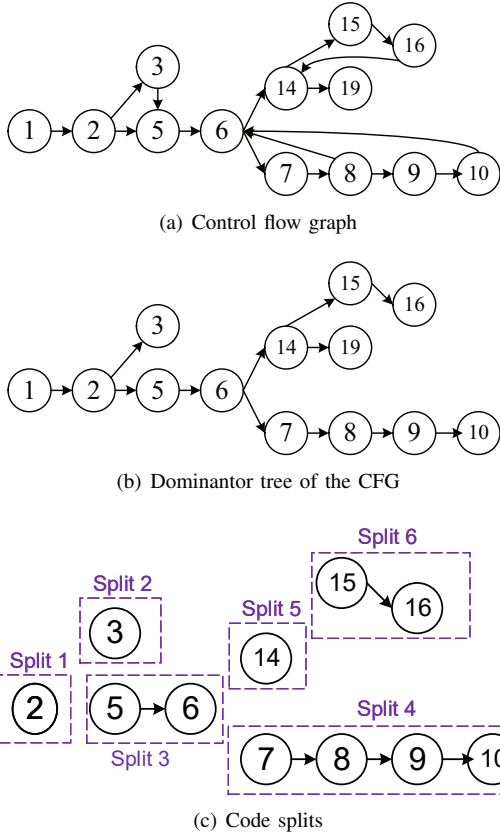


Fig. 3: Code splitting example

out-going edges. In the end, we will have multiple disconnected subgraphs, each corresponds to a set of statements in the code, i.e., a piece of split code.

- 4) We add the method declaration statement in the original complete code to the beginning of each split code, and use sophisticated tool to extract AST for each split code, i.e., *split AST*. Details of the AST generation tool will be described in Sec. VI-A.

Fig. 3 illustrates the steps of code splitting. The CFG for the example in Fig. 1 is shown in Fig. 3(a). The resulting dominator tree is in Fig. 3(b). Excluding the start and end node, the dominator tree is segmented to six parts, as shown in Fig. 3(c). We do not show the generated split AST for each split code due to page limits.

Note that, although there are various ways to split the code, we choose to leverage the blocks in the dominator tree [23] for two reasons:

- 1) The consecutive statements inside each block typically represent a piece of business logic which contains richer information than single statement (more analysis is provided in Sec. VI-C.).
- 2) There are no cyclic structures among blocks in the dominator tree and this greatly simplifies the design of the tree encoding phase (see the explanation of Eq. 3 in Sec. IV).

IV. TREE ENCODING

The tree encoding component targets at each split AST t , to learn a vectorized numerical representation, i.e., *syntax*

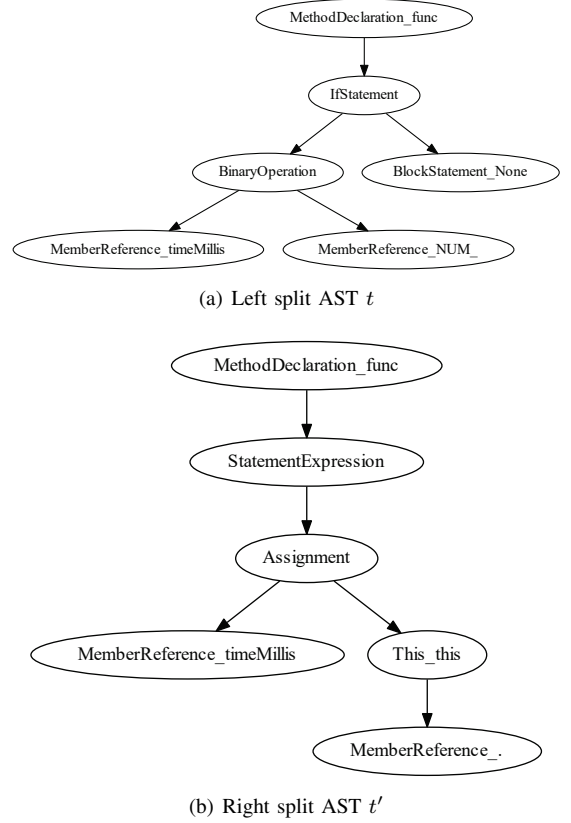


Fig. 4: An illustrative example of split ASTs

Algorithm 1: The procedure to obtain syntax embedding for split AST t

Input: AST t with node set $\mathcal{V}(t)$
Output: Syntax embedding \mathbf{e}_t

- 1 $v \leftarrow \text{root}(t)$;
- 2 $\mathbf{h}_v \leftarrow \text{Tree-LSTM}(v)$;
- 3 $\mathbf{e}_t \leftarrow \mathbf{h}_v$;
- 4 **Function** $\text{Tree-LSTM}(v)$:
 - 5 **if** $\mathcal{C}_v \neq \emptyset$ **then**
 - 6 **for** $c \in \mathcal{C}_v$ **do**
 - 7 $\mathbf{h}_c \leftarrow \text{Tree-LSTM}(c)$;
 - 8 Compute \mathbf{h}_v by Eq. 1;
 - 9 **else**
 - 10 Randomly initiate \mathbf{h}_c ;
 - 11 $\mathcal{C}_v \leftarrow \{c\}$;
 - 12 Compute \mathbf{h}_v by Eq. 1 ;
 - 13 **return** \mathbf{h}_v ;

embedding $\mathbf{e}_t \in \mathcal{R}^L$, where L is the embedding size. As mentioned in Sec. II, the syntax embedding is first pre-trained in a task-independent manner and then fine-tuned for the code summarization task. In this section, we introduce in detail the Syntax Embedding Pre-training (SEP) model for tree encoding in BASTS.

The SEP model takes two split ASTs t and t' as inputs, where an AST t has a set of nodes $\mathcal{V}(t)$. Each node $v \in \mathcal{V}(t)$ is assigned by a *type_value* in the AST generation step, which

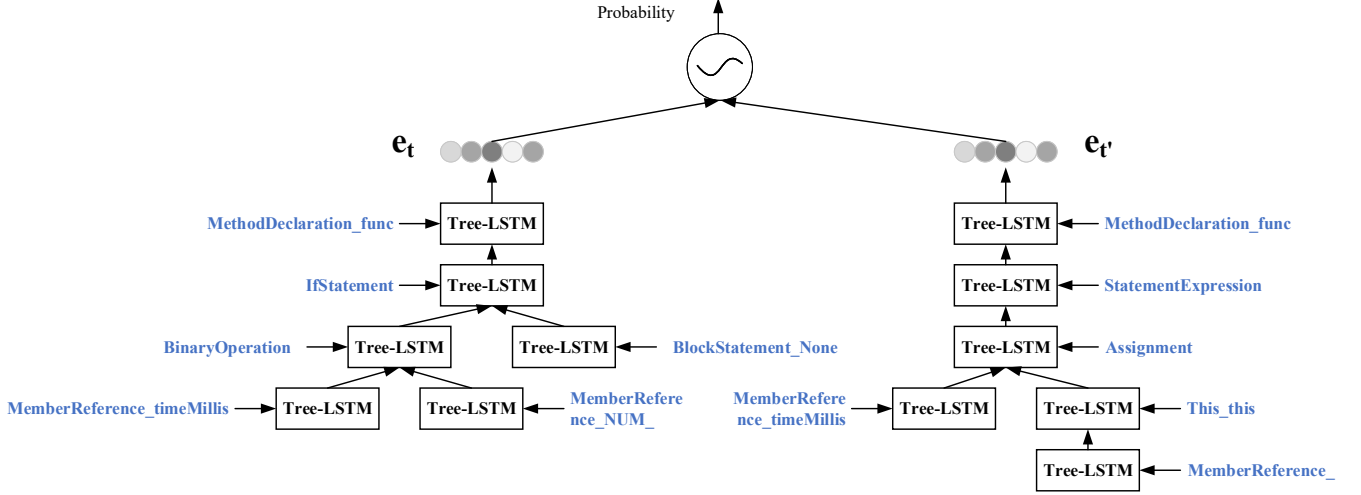


Fig. 5: An illustrative example of Syntax Embedding Pre-training (SEP) model on t, t'

is denoted as \mathbf{x}_v . For instance, the AST node with type `MemberReference` and value `timeMillis` will be assigned with `MemberReference_timeMillis`. Each non-leaf node v has a non-empty set of children nodes, which is defined as \mathcal{C}_v .

For each node v , we apply the Child-Sum Tree-LSTM [7] which outputs the hidden state $\mathbf{h}_v \in \mathcal{R}^L$ based on the inputs including hidden states of children nodes, i.e., \mathbf{h}_c ($c \in \mathcal{C}_v$) and the embedding vector of its type_value, i.e., \mathbf{x}_v . Given a non-leaf node v where $\mathcal{C}_v \neq \emptyset$, the Child-Sum Tree-LSTM operates as follows:

$$\begin{aligned}
 \tilde{\mathbf{h}}_v &= \sum_{c \in \mathcal{C}_v} \mathbf{h}_c \\
 \mathbf{i}_v &= \sigma(\mathbf{W}^i \mathbf{x}_v + \mathbf{U}^i \tilde{\mathbf{h}}_v + \mathbf{b}^i) \\
 \mathbf{f}_{v,c} &= \sigma(\mathbf{W}^f \mathbf{x}_v + \mathbf{U}^f \mathbf{h}_c + \mathbf{b}^f) \\
 \mathbf{o}_v &= \sigma(\mathbf{W}^o \mathbf{x}_v + \mathbf{U}^o \tilde{\mathbf{h}}_v + \mathbf{b}^o) \\
 \mathbf{u}_v &= \tanh(\mathbf{W}^u \mathbf{x}_v + \mathbf{U}^u \tilde{\mathbf{h}}_v + \mathbf{b}^u) \\
 \mathbf{m}_v &= \mathbf{i}_v \odot \mathbf{u}_v + \sum_{c \in \mathcal{C}_v} \mathbf{f}_{v,c} \odot \mathbf{m}_c \\
 \mathbf{h}_v &= \mathbf{o}_v \odot \tanh(\mathbf{m}_v)
 \end{aligned} \tag{1}$$

where \mathbf{i} , \mathbf{f} and \mathbf{o} are the *input*, *forget*, *output* gates, respectively. \mathbf{m} is the *memory cell*. \mathbf{b} is the bias vector, \mathbf{W} and \mathbf{U} are weight matrices, σ is the sigmoid activation function, \tanh is the tanh activation function, and \mathbf{x}_v is the input embedding of the type_value of node v .

For each split AST t , we use the output of the root node as the syntax embedding \mathbf{e}_t , which is computed according to Alg. 1. Note that $\tilde{\mathbf{h}}_v$ in Eq. 1 is the temporary hidden state summed over all children nodes $c \in \mathcal{C}_v$. For a leaf node v with $\mathcal{C}_v = \emptyset$, a “virtual” child node is randomly initialized to facilitate the computation of $\tilde{\mathbf{h}}_v$ (see lines 10-12 in Alg. 1). Finally, the SEP model predicts $f(t, t') \in (0, 1)$ based on the syntax embeddings \mathbf{e}_t and $\mathbf{e}_{t'}$:

$$f(t, t') = \sigma(\text{concat}(\mathbf{e}_t, \mathbf{e}_{t'})), \tag{2}$$

where σ is the sigmoid activation function, $\text{concat}(\cdot, \cdot) \in \mathcal{R}^{2L}$ is the concatenation of \mathbf{e}_t and $\mathbf{e}_{t'}$. In Fig. 4, we illustrate two split ASTs, which are some of the split ASTs after splitting the example in Fig. 1. Fig. 5 depicts the corresponding SEP model when learning these two split ASTs.

To learn the parameters, SEP optimizes the cross-entropy loss for next split AST prediction.

$$\mathcal{L}_{pre} = \sum_{t, t'} \left(\hat{f}(t, t') \log f(t, t') + (1 - \hat{f}(t, t')) \log (1 - f(t, t')) \right), \tag{3}$$

where $\hat{f}(t, t')$ is the gold standard label. If a split AST t is proceeded by t' in the dominator tree, i.e. $(t \rightarrow t') \in \mathcal{E}(DT)$, then $\hat{f}(t, t') = 1$, otherwise, $\hat{f}(t, t') = 0$. Note that \mathcal{L}_{pre} is the loss associated with source code, thus the pre-training is unsupervised and is less affected by the training data, e.g., the number of paired comments with source code and the quality of comments. It is worthy noting that our choice of code splitting based on dominator tree simplifies the design for predicting next split AST, since there is no cyclic structures after splitting. If the resulted code splits form a cycle graph (e.g., Split 1 is the subsequent block of Split 6 in Fig. 3(c)) after using other splitting methods, there will be an infinite number of (current split, next split) pairs. In this case, it will be necessary to define a strategy to limit the number of such pairs to be predicted in the pre-train task. Unfortunately, such a strategy is not easy to design, as any manual rule will make some pairs being learned less than others or hinder the information propagation along the cycle in the dependence learning. As a comparison, our approach explained in Sec. III opts to split based on dominator tree and eliminates the possibility of the cyclic structure.

V. SUMMARIZING CODE WITH TRANSFORMER

The overall code summarization model in BASTS, as shown in Fig. 6, follows an encoder-decoder framework [17], and we adopt Transformer [18] as the implementation.

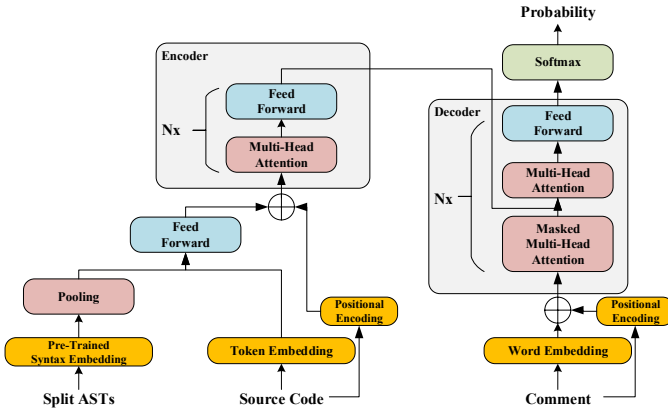


Fig. 6: Summarization model in BASTS

Input. Recall that the n -th training sample is defined as $\langle \mathbf{x}^n, \mathbf{y}^n \rangle$. In the sequel, the superscript n will be omitted when there is no ambiguity. Code splitting (Sec. III) and Syntax Embedding Pre-training (Sec. IV) obtain a set of L -dimensional vectors for the split AST, i.e., $\{\mathbf{e}_t | t \in \mathcal{T}\}$, where $\mathbf{e}_t \in \mathcal{R}^L$. The syntax embeddings for split ASTs flow through a pooling layer, which generates the representation of the overall syntax:

$$\text{AvgPooling}(\mathbf{e}_1, \dots, \mathbf{e}_{|\mathcal{T}|}) = \left[\frac{\sum_t \mathbf{e}_{t,1}}{|\mathcal{T}|}, \dots, \frac{\sum_t \mathbf{e}_{t,L}}{|\mathcal{T}|} \right]. \quad (4)$$

Each code token is also encoded by a numerical embedding vector. The syntax embeddings and code token embeddings are combined in a feed-forward layer. Note that each code embedding in a code sequence is concatenated with the same syntax embedding. Next, we will slightly abuse the notation here to use $\tilde{\mathbf{x}} \in \mathcal{R}^{2L}$ to represent the concatenation of the output of the pooling layer and one token embedding. $\tilde{\mathbf{x}}$ is fed into a single-layer feedforward neural network:

$$FFN(\tilde{\mathbf{x}}) = \text{ReLU}(\mathbf{W}^F \tilde{\mathbf{x}} + \mathbf{b}^F), \quad (5)$$

where ReLU is the ReLU activation function. \mathbf{W} and \mathbf{b} are weight matrix and vector, respectively. Intuitively, the output $FFN(\tilde{\mathbf{x}})$ is a weighted combination of syntax embedding and token embedding. This layer reveals that the understanding of source code incorporates both the syntax structure and the token sequence.

To keep the relative position of a token in the code, position encodings are added to the input embeddings. The position encoding is formatted by:

$$\text{pos}(d, l) = \begin{cases} \sin(d/10000^{l/L}) & \text{if } l \text{ is even} \\ \cos(d/10000^{l/L}) & \text{if } l \text{ is odd} \end{cases} \quad (6)$$

where d is the token index in the source code, l is the index of the embedding vector, and L is the embedding size. \sin and \cos are the sin and cosine functions, respectively.

Encoder. The encoder is responsible for reading the input one token at a time and encoding the current information in a hidden state. The input embedding vector flows through a stack of encoder layers, where each encoder layer consists of a multi-head self-attention layer and a feedforward layer.

The multi-head self-attention layer first constructs the input matrix $\mathbf{X} \in \mathcal{R}^{D \times L}$ by stacking D_n tokens and padding the rest $D - D_n$ rows with zeros. Three $L \times L$ trainable matrices \mathbf{W}^Q , \mathbf{W}^K and \mathbf{W}^V transform the input matrix to *queries* $\mathbf{Q} = \mathbf{X}\mathbf{W}^Q$, *keys* $\mathbf{K} = \mathbf{X}\mathbf{W}^K$ and *values* $\mathbf{V} = \mathbf{X}\mathbf{W}^V$. The multi-head mechanism splits the *queries*, *keys* and *values* to H parts, each of \mathbf{Q}^h , \mathbf{K}^h and \mathbf{V}^h has the dimension size $D \times L_h$ so that each head can focus on different parts of the input. The output of each head is computed as:

$$\text{Attention}(\mathbf{Q}^h, \mathbf{K}^h, \mathbf{V}^h) = \text{softmax}\left(\frac{\mathbf{Q}^h \mathbf{K}^{hT}}{\sqrt{L_h}}\right) \mathbf{V}^h. \quad (7)$$

The outputs of each head are concatenated and fed to a single-layer feedforward neural network. After that, each encoder layer also employs residual connection and normalization [18].

Decoder. The decoder is responsible for generating a comment word based on the previous output and the current encoding state. The decoder is composed of a stack of identical layers. In each layer, there are three sublayers. The first sublayer is a masked multi-head self-attention mechanism. The second sublayer is a multi-head self-attention layer. The operation of the multi-head attention in the decoder is similar to Eq. 7. The difference between masked multi-head self-attention and multi-head self-attention is that the former looks to earlier positions in the output. The third sublayer is a fully connected feedforward network. Finally, the output of the stacked layers is passed through a softmax layer to produce the probabilities of each candidate word.

Optimization. The probabilities output by the decoder will be used in the standard cross-entropy loss which can be minimized by algorithms such as gradient descent to estimate the parameters of the overall summarization model.

VI. EXPERIMENT

We conduct experiments to analyze the performance of BASTS¹. We focus on the following research questions:

RQ1: Does BASTS generate better comments, compared with the state-of-the-art approaches?

RQ2: Does each component contribute to the performance of BASTS?

RQ3: What are the impacts of dataset characteristics on the quality of generated comments?

A. Experimental Setup

Data. We use two public benchmarks: the Java dataset² [12] which was collected from GitHub's Java repositories, and the Python dataset³ [26] which is provided by the CodeSearchNet challenge. We replace number, string and boolean variables with special tokens $\langle \text{NUM} \rangle$, $\langle \text{STR} \rangle$ and $\langle \text{BOOL} \rangle$ in source code, respectively. For both datasets, we use SCITOOLS⁴ to

¹Our implementation is available at <https://github.com/XMUDM/BASTS>.

²<https://github.com/xing-hu/EMSE-DeepCom>

³<https://github.com/github/CodeSearchNet>

⁴<https://scitools.com/support/modifying-the-control-flow-graph>

generate CFGs. We also use the Javalang library⁵ and the ast library⁶ to generate ASTs for Java dataset and Python dataset, respectively. The source code in Java is parsed into tokens by Javalang. Python dataset has been tokenized already by its authors. For both datasets, identifiers are split into subtokens (i.e., a variable `camelCase` yields two subtokens `camel` and `case`), and comments are parsed into words by NLTK⁷.

The statistics of datasets are shown in Tab. 1. For both datasets, we use the default 90/5/5 splits provided by the authors [12], [26] for training, validation and test sets. It is worth pointing out that, for Java dataset, there are 5,824 instance in the validation set and 5,341 instances in the test set which also exist in the training set. We remove them from the validation/test set. We also delete erroneous methods which can not generate CFGs or ASTs in two datasets.

Hardware. The experiments were run on a machine with two Intel(R) Xeon(R) CPU E5-2678 v3 @ 2.50GHz, 256 GB main memory and a GeForce RTX 2080 Ti graphics card.

Evaluation Metrics. We use BLEU, METEOR and ROUGE which are widely used in previous works for code summarization [27]. BLEU basically computes the averaged percentage of n-gram matches over each $n = 1, \dots, N$:

$$P_n = \frac{\sum_{gram_n \in s} \#match(gram_n)}{\#(gram_n)}, \quad BLEU = \rho \left(\prod_n P_n \right)^{1/N}$$

where $\sum_{gram_n \in s} \#match(gram_n)$ is the number of matching n-grams in the result and the gold standard s , $\#(gram_n)$ is the number of n-grams in the result, ρ is a brevity penalty. We adopt S-BLEU and C-BLEU, which indicate average sentence-level BLEU score and corpus-level average BLEU score, respectively.

ROUGE compares the generated comment with the gold-standard comment by counting the number of their overlapping textual units. We consider textual units such as unigram, bigram and longest common sequence (LCS), i.e., ROUGE-1, ROUGE-2 and ROUGE-L. We adopt F-Score for detailed calculation. For example, ROUGE-L F-Score is computed as:

$$R = \frac{\#LCS(s, y)}{m}, \quad P = \frac{\#LCS(s, y)}{n}, \quad F = \frac{2RP}{R + P}$$

where $\#LCS(s, y)$ is the number of matching longest common sequence between the gold-standard and the result summary, m is the number of LCSs in the result comment and n is the number of LCSs in the gold standard summary. The computation of ROUGE-1 and ROUGE-2 is similar by replacing LCS with matching unigrams or bigrams.

METEOR⁸ scores the quality of translation by aligning them to one or more gold-standards. Alignments are based on exact, stem, synonym, and paraphrase matches between words and phrases. The metric includes several free parameters that are tuned to emulate human judgments.

⁵<https://pypi.org/project/javalang>

⁶<https://docs.python.org/3/library/ast.html>

⁷<https://www.nltk.org>

⁸<http://www.cs.cmu.edu/~alavie/METEOR/>

Table 1: Statistics of datasets

Dataset	#Methods	#Tokens	#Words	#Split ASTs
Java	441,487	45,790	55,015	1,273,301
Python	241,322	237,939	75,883	1,033,210

B. Code Summarization Performance (RQ1)

To answer RQ1, we compare BASTS with seven state-of-the-art methods:

- 1) **CODE-NN**⁹ [19] uses LSTM and a global attention mechanism in the seq2seq framework.
- 2) **Hybrid-DRL**¹⁰ [20] applies an actor-critic network. The actor network produces the hidden space. The critic network adopts BLEU as the reward function.
- 3) **Hybrid-DeepCom**¹¹ [12] uses GRU to build a code encoder for code tokens, and an AST encoder for serialized ASTs from SBT. A hybrid attention component fuses the lexical and syntactical information.
- 4) **AST-attendgru**¹² [11] applies two GRU encoders on code sequences and serialized ASTs from SBT. A different attention mechanism compared to Hybrid-DeepCom is used for information fusion.
- 5) **ASTNN**¹³ [21] splits an AST into a sequence of statement trees. Statement trees are encoded by an RvNN based encoder. The learned sequence of statement tree representations are fed into GRU to generate the representation of the original code for code classification and clone detection. For a fair comparison, we feed the code representation learned by ASTNN into Eq. 4 for code summarization, and other components are the same as BASTS.
- 6) **Transformer**¹⁴ [28] employs the Transformer model [18] in code summarization task.
- 7) **Rencos**¹⁵ [13] retrieves two code snippets that are most similar to the input code snippet. Then, it encodes the input and two retrieved code snippets, and generates the summary by fusing them during decoding.

We use the default hyper-parameter settings provided by each method. Following Hu et al. [12], we set maximal code length to 100, maximal comment length to 30, and maximal epochs to 50 for all methods (we also report the impact of different code/comment lengths in Sec. VI-D). We use 512 as the embedding size and 128 as the batch size. The exception is that we use the batch size 64 for ASTNN since ASTNN with the batch size 128 will encounter out-of-memory error. For a fair comparison, we also compare ASTNN with BASTS when both methods are trained with the batch size 64 in Sec. VI-C. We use 8 heads in Transformer and BASTS. We adopt greedy search for AST-attendgru and Hybrid-DRL, following their original design. We use beam search with beam size 5 for CODE-NN and Hybrid-DeepCom, as mentioned in

⁹<https://github.com/sriniyer/codenn>

¹⁰https://github.com/wanyao1992/code_summarization_public

¹¹<https://github.com/xing-hu/EMSE-DeepCom>

¹²<https://bit.ly/2MLSxFg>

¹³<https://github.com/zhangj111/astnn>

¹⁴<https://github.com/wasiahmad/NeuralCodeSum>

¹⁵<https://github.com/zhangj111/rencos>

Table 2: Comparative code summarization performance on Java dataset. Best performance are in bold font, and improvement percentage of BASTS w.r.t the best baseline are in brackets.

Method	S-BLEU	C-BLEU	METEOR	ROUGE-1	ROUGE-2	ROUGE-L
CODE-NN	24.22	14.68	14.79	34.60	17.85	33.18
Hybrid-DRL	20.16	10.99	13.15	30.45	13.80	29.22
Hybrid-DeepCom	29.19	19.82	18.21	40.52	23.46	38.70
AST-attendgru	27.50	17.85	16.44	36.73	20.83	35.32
ASTNN (batch size=64)	25.40	15.94	16.92	38.21	19.66	36.10
Transformer	30.98	23.63	20.30	43.49	27.60	42.20
Rencos	32.90	26.03	21.52	44.12	29.31	42.84
BASTS	36.38 (↑ 10.58%)	30.37 (↑ 16.67%)	24.12 (↑ 12.08%)	49.63 (↑ 12.49%)	33.13 (↑ 13.03%)	47.85 (↑ 11.69%)

Table 3: Comparative code summarization performance on Python dataset. Best performance are in bold font, and improvement percentage of BASTS w.r.t the best baseline are in brackets.

Method	S-BLEU	C-BLEU	METEOR	ROUGE-1	ROUGE-2	ROUGE-L
CODE-NN	15.37	1.68	7.25	18.57	3.28	17.02
Hybrid-DRL	14.09	2.25	7.80	18.96	4.18	17.77
Hybrid-DeepCom	10.24	0.85	5.03	13.35	2.36	12.41
AST-attendgru	14.47	1.17	6.42	17.05	2.85	15.76
Transformer	14.80	1.64	7.60	19.30	4.00	17.91
Rencos	15.03	3.31	8.73	20.90	5.51	19.40
BASTS	15.97 (↑ 6.25%)	2.77 (↓ 16.31%)	9.79 (↑ 12.14%)	24.61 (↑ 17.75%)	6.18 (↑ 12.16%)	22.72 (↑ 17.11%)

their papers. For BASTS, greedy search is used for simplicity. We adopt adam [29] for optimization.

Tab. 2 and Tab. 3 report the results on two datasets. Note that the authors of ASTNN do not provide the implementation for splitting python code using their method. Hence we only evaluate ASTNN on Java dataset. From the results, we can observe that BASTS is robustly better than state-of-the-art competitors. In most cases, BASTS exceeds baselines by more than 10% in terms of all measures. As pointed out by Gros et al. [27], existing neural code summarization methods may not indeed surpass traditional retrieval-based methods. Therefore, we compare BASTS with Rencos which combines the merits of both retrieval-based and neural network-based methods. From Tab. 2 and Tab. 3, we can see that BASTS has better performance than Rencos in almost all cases, showing that BASTS is better than the retrieval-based model.

C. Contribution of Different Components of BASTS (RQ2)

To answer RQ2, we further investigate Transformer, ASTNN and BASTS, which share some components in their models. We report the performance of BASTS_{b128} (batch size=128, i.e., BASTS), BASTS_{b64} (batch size=64), Transformer (batch size=128) and ASTNN (batch size=64) at each epoch ($i = 1, \dots, 100$) on Java dataset in Fig. 7. Due to page limit, we omit the similar results on Python dataset.

Recently, the Transformer model has achieved a great success in various domains [30]. To verify whether the superiority of BASTS over existing methods is mainly caused by the use of the Transformer architecture, we compare BASTS with the baseline Transformer [28] which adopts the pure Transformer architecture for code summarization in Tab. 2, Tab. 3 and Fig. 7. From the results, we can see that BASTS has much better performance than the pure Transformer architecture in code summarization task, showing that components like

code splitting and pre-training in BASTS indeed help improve its performance. Thus, we can conclude that using a more advanced neural architecture is not the only reason that BASTS outperforms existing methods.

The next question is whether the way that BASTS splits ASTs is better than the existing AST splitting method, i.e., split ASTs into statement trees like ASTNN. Recall that the difference between ASTNN and BASTS in our experiments is how they split ASTs, and other components in these two methods are the same. In the main experiments in Tab. 2 and Tab. 3, the results of ASTNN are obtained when it uses the batch size 64 to avoid the out-of-memory error. If we compare ASTNN and BASTS_{b64} which both use the batch size 64 in Fig. 7, we can observe that, after around 10 epochs, BASTS_{b64} starts to significantly outperform ASTNN. We further conduct case studies on the generated comments from ASTNN and BASTS, and the statements decomposed by ASTNN. We show one case study from Java dataset in Fig. 10. From Fig. 10, we can see that the inadequate summary generated by ASTNN is related to the first statement only. The reason is that each statement tree split by ASTNN only consists of one statement, and ASTNN cannot capture the information beyond the statement boundary using its RvNN encoder only. If the sequential neural model (ASTNN uses GRU) which is used to model the sequence of statements does not fuse statement representations well, the generated comment will only include local information from one statement. BASTS generates a perfect comment in Fig. 10. Assume that BASTS also encounters the problem that its sequential model does not fuse the representations of split ASTs well like ASTNN in Fig. 10. Then, for such cases, BASTS will still generate a comment containing more information than a single statement since split ASTs in BASTS is based on blocks in the dominator tree (one block consists of several statements). From above

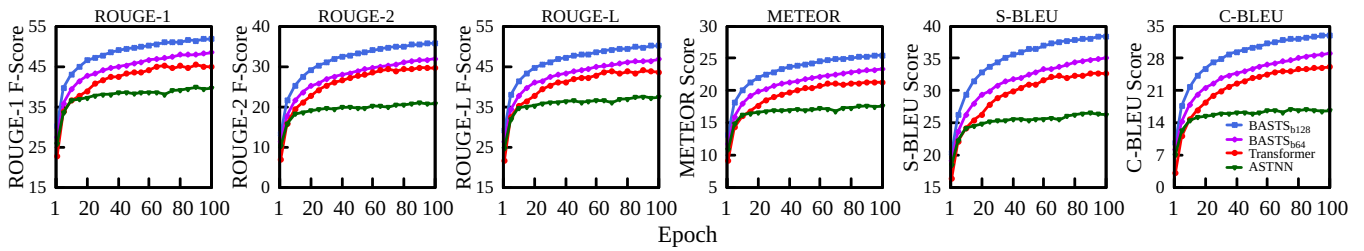


Fig. 7: Performance of different methods at each epoch on Java dataset.

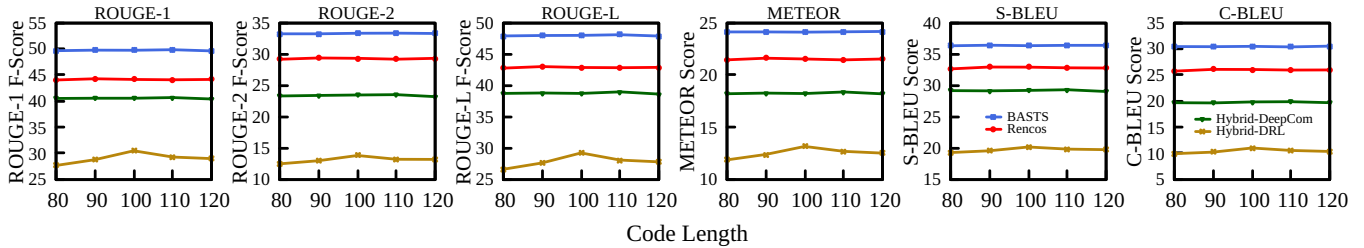


Fig. 8: Performance w.r.t varying code length on Java dataset.

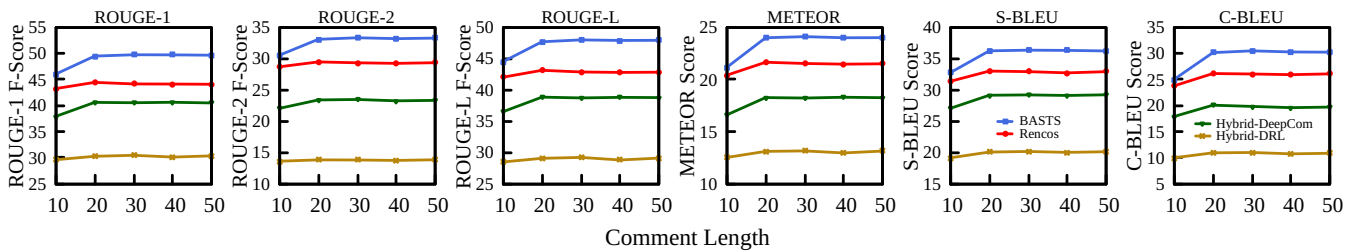


Fig. 9: Performance w.r.t varying comment length on Java dataset.

```

1 public static void writeDescriptor(Descriptor descriptor, File file,
2   String encoding, boolean isIndent) throws IOException {
3   OutputStream out = null;
4
5   try {
6     out = new FileOutputStream(file);
7     writeDescriptor(descriptor, out, encoding, isIndent);
8   } finally {
9     if (out != null) {
10      try {
11        out.close();
12      } catch (IOException ioe) {}
13    }
14  }
15 }
16 }

```

Reference: writes the specified document to a file
 BASTS: writes the specified document to a file
 ASTNN: writes a boolean value

Fig. 10: Case Study: A comparison between ASTNN and BASTS for test sample #3797 in Java dataset. Red dashed lines indicate the statements decomposed by ASTNN.

observations and analysis, we can conclude that splitting ASTs based on the blocks in the dominator tree like BASTS is better than splitting ASTs by statements like ASTNN.

D. Impact of Dataset Characteristics (RQ3)

Following Wan et al. [20], we analyze the impact of code length and comment length of the Java data on the quality

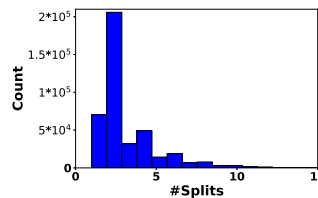


Fig. 11: Distribution of the number of split ASTs in Java dataset.

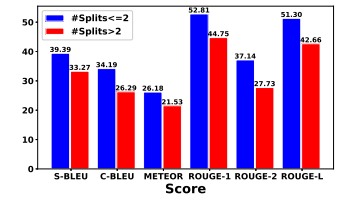


Fig. 12: Results on two partitions with different number of split ASTs in Java dataset.

of generated comments since the code length may affect the representation learning and the comment length may affect the evaluation of comment generation. Due to space limit, we do not show the similar analysis of Python dataset.

As shown in Fig. 8 and Fig. 9, BASTS has best performance w.r.t. varying values of code and comment length on Java dataset, compared to baselines Rencos, Hybrid-DeepCom and Hybrid-DRL. From Fig. 8 and Fig. 9, we can see that BASTS performs best when code length is 100 or comment length is 30. The underlying reason is that in Java dataset, 83% of the methods are with code length ≤ 100 , and 99.99% of the comments are with length ≤ 30 .

We next visualize the distribution of split ASTs by BASTS in Java dataset in Fig. 11. We observe that most methods can be segmented to two splits. To investigate whether the number

of split ASTs affects BASTS’s performance, we divide the Java dataset into two partitions, one consisting of instances that can be segmented to at most two splits and the other consisting of instances that can be segmented to more than two splits. Note that the partitioning is performed on both training and test sets. Fig. 12 reports the performance of BASTS on the two partitions. We can see that BASTS performs slightly worse on the partition with more splits, showing that it is more difficult to summarize such complex code (with more split ASTs). Observing both Fig. 12 and Tab. 2, we can see that training BASTS with complex code only (i.e., partition with $\#splits > 2$) still give better performance than Rencos (which is the best competitor as shown in Tab. 2) trained with all code in 4 out of 6 evaluation metrics, which further demonstrates the superiority of BASTS over existing methods.

VII. RELATED WORK

A. Source Code Summarization

Source code comments are important for facilitating collaboration between software developers. Much effort has been devoted to designing automatic code summarization techniques in the past decade. Conventional approaches mostly adopt Information Retrieval (IR) techniques [31]–[42] or methods from NLP domain [43]–[47].

Recent breakthroughs of deep learning has significantly affect the design of automatic code summarization methods [3]. Most neural code summarization methods leverage the seq2seq framework in machine translation [17]. Early approaches leverage various deep learning techniques like LSTM [19], convolutional attention network [48], dual training [49] or multiple encoders [50] to generate code summarization based on input code tokens and the seq2seq framework.

Recently, there is a surge of works on using ASTs of the source code to guide the generation of summarization. Most of them firstly serialize ASTs into sequences. One representative traversal method is the Structure-Based Traversal method (SBT) [4]. Given the serialized ASTs, existing works adopt sequential neural networks like LSTM [4], GRU [11], [12] or Transformer [28] to extract features for generating code summarization. The limitation of these methods is that structural information cannot be fully captured from serialized ASTs. To overcome this problem, some works consider replacing the RNN based encoder with a Graph Neural Network (GNN) based encoder [10] or feeding the outputs from the sequential encoder as the initial node representations into the GNN based encoder [51]. These methods typically encounter high overhead and face difficulties in training, as ASTs are deep with many nodes which are hard and slow to learn. In order to balance the retention of structural information and the training difficulties, Wan et al. [20] transform ASTs into binary trees for efficient training at the cost of partial information loss. Besides above methods of modeling ASTs, another way is to represent a code snippet as the set of compositional paths in its AST for code summarization [5].

There are other directions for improving code summarization. Wan et al. [20] argue that the encoder-decoder

framework suffers from exposure bias, and they opt to deep reinforcement learning for better code summarization. Chen and Zhou [52] consider modeling code retrieval and code summarization tasks together. Zhang et al. [13] design a new method which combines merits of both IR-based and neural network-based code summarization methods. Haque et al. [53] add an additional RNN encoder to the seq2seq framework to encode file context and enhance code summarization. Gros et al. [27] reviews the evaluation metrics used in existing code summarization works.

B. Representation Learning for Programs

The large volume of public available source code in source code management platforms (e.g., GitHub) makes it possible to extract and learn knowledge from code corpora. Hence, a tremendous amount of works on enhancing program representation learning and comprehension emerges recently [1], [2].

Some recent works also study modeling ASTs and scale training to large programs. Allamanis et al. [54] represent programs as graphs. Gated Graph Neural Network [55] is adopted to learn the representation of programs for variable misuse and variable naming prediction. Alon et al. [56] decompose an AST as a collection of paths and generate the representations of code by learning how to aggregate paths. Wang and Li [57] tailor GNN to modeling AST for code completion task. Zhang et al. [21] split an AST into statements and use a preorder traversal to transform the AST to a sequence of statements. Then, the statement sequence is fed into RvNN to construct the statement representations which are finally passed to GRU for code classification and clone detection. Although their idea resembles our work, they decompose ASTs based on individual statements. This way, the information beyond the boundary of statements is hard to capture.

Inspired by the success of the pre-trained model BERT [25], pre-training has also been applied in program representation learning. CodeBERT [58] is a bimodal pre-trained model on programming language and natural language, and it is designed to benefit down-stream tasks. However, CodeBERT produces poorer results compared to methods which make use of AST paths [58] since it only takes original code as the input.

VIII. CONCLUSION

In this paper, we have presented a novel method BASTS to improve code summarization. To alleviate the problem of difficult and slow training on ASTs, BASTS uses block-wise code splitting and localizes the hierarchical syntax structure of AST generated for each code split. BASTS further adopts a pre-training strategy to capture local non-linear syntax encoding of split ASTs, which helps its Transformer-based summarization model generate high-quality code summaries. Experiments on public benchmarks have demonstrated that BASTS significantly outperform state-of-the-art models. In the future, we plan to exploit more sophisticated program slicing techniques and combine code slicing based on variables with variable-aware neural summarization models to further improve BASTS.

ACKNOWLEDGMENT

Chen Lin is supported by the National Natural Science Foundation of China (No. 61972328), and Joint Innovation Research Program of Fujian Province of China (No. 2020R0130). Hui Li is supported by the National Natural Science Foundation of China (No. 62002303, 62036004), and the Natural Science Foundation of Fujian Province of China (No. 2020J05001). Rongxin Wu is supported by the National Natural Science Foundation of China (No. 61902329).

REFERENCES

- [1] M. Allamanis, E. T. Barr, P. T. Devanbu, and C. A. Sutton, "A survey of machine learning for big code and naturalness," *ACM Comput. Surv.*, vol. 51, no. 4, pp. 81:1–81:37, 2018.
- [2] T. H. M. Le, H. Chen, and M. A. Babar, "Deep learning for source code modeling and generation: Models, applications, and challenges," *ACM Comput. Surv.*, vol. 53, no. 3, pp. 62:1–62:38, 2020.
- [3] Y. Zhu and M. Pan, "Automatic code summarization: A systematic literature review," *arXiv Preprint*, 2019. [Online]. Available: <https://arxiv.org/abs/1909.04352>
- [4] X. Hu, G. Li, X. Xia, D. Lo, and Z. Jin, "Deep code comment generation," in *ICPC*, 2018, pp. 200–210.
- [5] U. Alon, S. Brody, O. Levy, and E. Yahav, "code2seq: Generating sequences from structured representations of code," in *ICLR*, 2019.
- [6] V. Kovalenko, E. Bogomolov, T. Bryksin, and A. Bacchelli, "Pathminer: a library for mining of path-based representations of code," in *MSR*, 2019, pp. 13–17.
- [7] K. S. Tai, R. Socher, and C. D. Manning, "Improved semantic representations from tree-structured long short-term memory networks," in *ACL*, vol. 1, 2015, pp. 1556–1566.
- [8] X. Zhu, P. Sobhani, and H. Guo, "Long short-term memory over recursive structures," in *ICML*, vol. 37, 2015, pp. 1604–1612.
- [9] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," *IEEE Trans. Knowl. Data Eng.*, 2020.
- [10] A. LeClair, S. Haque, L. Wu, and C. McMillan, "Improved code summarization via a graph neural network," in *ICPC*, 2020, pp. 184–195.
- [11] A. LeClair, S. Jiang, and C. McMillan, "A neural model for generating natural language summaries of program subroutines," in *ICSE*, 2019, pp. 795–806.
- [12] X. Hu, G. Li, X. Xia, D. Lo, and Z. Jin, "Deep code comment generation with hybrid lexical and syntactical information," *Empirical Software Engineering*, 2019.
- [13] J. Zhang, X. Wang, H. Zhang, H. Sun, and X. Liu, "Retrieval-based neural source code summarization," in *ICSE*, 2020, pp. 1385–1397.
- [14] T. Young, D. Hazarika, S. Poria, and E. Cambria, "Recent trends in deep learning based natural language processing [review article]," *IEEE Comput. Intell. Mag.*, vol. 13, no. 3, pp. 55–75, 2018.
- [15] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, no. 6088, pp. 533–536, 1986.
- [16] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [17] K. Cho, B. van Merriënboer, Ç. Gülçehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," in *EMNLP*, 2014, pp. 1724–1734.
- [18] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *NIPS*, 2017, pp. 5998–6008.
- [19] S. Iyer, I. Konstas, A. Cheung, and L. Zettlemoyer, "Summarizing source code using a neural attention model," in *ACL*, vol. 1, 2016, pp. 2073–2083.
- [20] Y. Wan, Z. Zhao, M. Yang, G. Xu, H. Ying, J. Wu, and P. S. Yu, "Improving automatic source code summarization via deep reinforcement learning," in *ASE*, 2018, pp. 397–407.
- [21] J. Zhang, X. Wang, H. Zhang, H. Sun, K. Wang, and X. Liu, "A novel neural source code representation based on abstract syntax tree," in *ICSE*, 2019, pp. 783–794.
- [22] R. Socher, C. C. Lin, A. Y. Ng, and C. D. Manning, "Parsing natural scenes and natural language with recursive neural networks," in *ICML*, 2011, pp. 129–136.
- [23] R. T. Prosser, "Applications of boolean matrices to the analysis of flow diagrams," in *IRE-AIEE-ACM*, 1959, pp. 133–138.
- [24] F. E. Allen, "Control flow analysis," *ACM Sigplan Notices*, vol. 5, no. 7, pp. 1–19, 1970.
- [25] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in *NAACL-HLT*, vol. 1, 2019, pp. 4171–4186.
- [26] H. Husain, H. Wu, T. Gazit, M. Allamanis, and M. Brockschmidt, "Codesearchnet challenge: Evaluating the state of semantic code search," *arXiv Preprint*, 2019. [Online]. Available: <https://arxiv.org/abs/1909.09436>
- [27] D. Gros, H. Sezhiyan, P. Devanbu, and Z. Yu, "Code to comment "Translation": Data, metrics, baselining & evaluation," in *ASE*, 2020, pp. 746–757.
- [28] W. U. Ahmad, S. Chakraborty, B. Ray, and K. Chang, "A transformer-based approach for source code summarization," in *ACL*, 2020, pp. 4998–5007.
- [29] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICLR*, 2015.
- [30] Y. Tay, M. Dehghani, D. Bahri, and D. Metzler, "Efficient transformers: A survey," *arXiv Preprint*, 2020. [Online]. Available: <https://arxiv.org/abs/2009.06732>
- [31] S. Haiduc, J. Aponte, L. Moreno, and A. Marcus, "On the use of automated text summarization techniques for summarizing source code," in *WCRE*, 2010, pp. 35–44.
- [32] G. Salton, A. Wong, and C. Yang, "A vector space model for automatic indexing," *Commun. ACM*, vol. 18, no. 11, pp. 613–620, 1975.
- [33] S. C. Deerwester, S. T. Dumais, T. K. Landauer, G. W. Furnas, and R. A. Harshman, "Indexing by latent semantic analysis," *JASIS*, vol. 41, no. 6, pp. 391–407, 1990.
- [34] S. Haiduc, J. Aponte, and A. Marcus, "Supporting program comprehension with source code summarization," in *ICSE*, vol. 2, 2010, pp. 223–226.
- [35] P. Rodeghero, C. McMillan, P. W. McBurney, N. Bosch, and S. K. D’Mello, "Improving automated source code summarization via an eye-tracking study of programmers," in *ICSE*, 2014, pp. 390–401.
- [36] A. N. Aizawa, "An information-theoretic perspective of tf-idf measures," *Inf. Process. Manag.*, vol. 39, no. 1, pp. 45–65, 2003.
- [37] B. P. Eddy, J. A. Robinson, N. A. Kraft, and J. C. Carver, "Evaluating source code summarization techniques: Replication and expansion," in *ICPC*, 2013, pp. 13–22.
- [38] D. M. Mimno, W. Li, and A. McCallum, "Mixtures of hierarchical topics with pachinko allocation," in *ICML*, vol. 227, 2007, pp. 633–640.
- [39] D. Movshovitz-Attias and W. W. Cohen, "Natural language models for predicting programming comments," in *ACL*, vol. 2, 2013, pp. 35–40.
- [40] J. M. Fowkes, P. Chanthirasegaran, R. Ranca, M. Allamanis, M. Lapata, and C. Sutton, "TASSAL: autofolding for source code summarization," in *ICSE (Companion Volume)*, 2016, pp. 649–652.
- [41] J. M. Fowkes, P. Chanthirasegaran, R. Ranca, M. Allamanis, M. Lapata, and C. Sutton, "Autofolding for source code summarization," *IEEE Trans. Software Eng.*, vol. 43, no. 12, pp. 1095–1109, 2017.
- [42] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," in *NIPS*, 2001, pp. 601–608.
- [43] G. Sridhara, E. Hill, D. Muppaneni, L. L. Pollock, and K. Vijay-Shanker, "Towards automatically generating summary comments for java methods," in *ASE*, 2010, pp. 43–52.
- [44] G. Sridhara, L. L. Pollock, and K. Vijay-Shanker, "Generating parameter comments and integrating with method summaries," in *ICPC*, 2011, pp. 71–80.
- [45] G. Sridhara, L. L. Pollock, and K. Vijay-Shanker, "Automatically detecting and describing high level actions within methods," in *ICSE*, 2011, pp. 101–110.
- [46] P. W. McBurney and C. McMillan, "Automatic documentation generation via source code summarization of method context," in *ICPC*, 2014, pp. 279–290.
- [47] X. Wang, L. L. Pollock, and K. Vijay-Shanker, "Automatically generating natural language descriptions for object-related statement sequences," in *SANER*, 2017, pp. 205–216.
- [48] M. Allamanis, H. Peng, and C. Sutton, "A convolutional attention network for extreme summarization of source code," in *ICML*, vol. 48, 2016, pp. 2091–2100.

- [49] B. Wei, G. Li, X. Xia, Z. Fu, and Z. Jin, "Code generation as a dual task of code summarization," in *NeurIPS*, 2019, pp. 6559–6569.
- [50] X. Hu, G. Li, X. Xia, D. Lo, S. Lu, and Z. Jin, "Summarizing source code with transferred API knowledge," in *IJCAI*, 2018, pp. 2269–2275.
- [51] P. Fernandes, M. Allamanis, and M. Brockschmidt, "Structured neural summarization," in *ICLR*, 2019.
- [52] Q. Chen and M. Zhou, "A neural framework for retrieval and summarization of source code," in *ASE*, 2018, pp. 826–831.
- [53] S. Haque, A. LeClair, L. Wu, and C. McMillan, "Improved automatic summarization of subroutines via attention to file context," in *MSR*, 2020, pp. 300–310.
- [54] M. Allamanis, M. Brockschmidt, and M. Khademi, "Learning to represent programs with graphs," in *ICLR*, 2018.
- [55] Y. Li, D. Tarlow, M. Brockschmidt, and R. S. Zemel, "Gated graph sequence neural networks," in *ICLR*, 2016.
- [56] U. Alon, M. Zilberstein, O. Levy, and E. Yahav, "code2vec: learning distributed representations of code," *Proc. ACM Program. Lang.*, vol. 3, no. POPL, pp. 40:1–40:29, 2019.
- [57] Y. Wang and H. Li, "Code completion by modeling flattened abstract syntax trees as graphs," *arXiv Preprint*, 2021. [Online]. Available: <https://arxiv.org/abs/2103.09499>
- [58] Z. Feng, D. Guo, D. Tang, N. Duan, X. Feng, M. Gong, L. Shou, B. Qin, T. Liu, D. Jiang, and M. Zhou, "Codebert: A pre-trained model for programming and natural languages," pp. 1536–1547, 2020.