

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Improving Automatic Source Code Summarization via Deep Reinforcement Learning

Abstract—Code summarization provides a high level natural language description of the function performed by code, as it can benefit the software maintenance, code categorization and retrieval. To the best of our knowledge, most state-of-the-art approaches follow an encoder-decoder framework which encodes the code into a hidden space and then decodes it into natural language space, suffering from two major drawbacks: a) their encoders only consider the sequential content of code, ignoring the tree structure which is also critical for the task of code summarization; b) their decoders are typically trained to predict the next word by maximizing the likelihood of next ground-truth word with previous ground-truth word given. However, it is expected to generate the entire sequence from scratch at test time. This discrepancy can cause an *exposure bias* issue, making the learnt decoder suboptimal. In this paper, we incorporate an abstract syntax tree structure as well as sequential content of code snippets into a deep reinforcement learning framework (i.e., actor-critic network learning). The actor network provides the confidence of predicting the next word according to current state. On the other hand, the critic network evaluates the reward value of all possible extensions of the current state and can provide global and lookahead guidance for explorations. We employ an advantage reward composed of BLEU metric to train both networks. Comprehensive experiments on a real-world dataset show the effectiveness of our proposed model when compared with the state-of-the-art ones.

I. INTRODUCTION

In the life cycle of software development (e.g., implementation, testing and maintenance), nearly 90% of effort is used for maintenance, and much of this efforts are spent on understanding the maintenance task and related software source codes [1]. Thus, documentation which provides a high level description of the task performed by code is always a must for software maintenance. Even though various techniques have been developed to facilitate the programmer during the implementation and testing of software, documenting code with comments remains a labour-intensive task, making few real-world software projects adequately document the code to reduce future maintenance costs [2], [3].

It’s nontrivial for a novice programmer to write good comments for source code. A good comment should at least have the following characteristics: a) Correctness. The comments should correctly clarify the intent of code. b) Fluency. The comments should be fluent natural language that maintainers can easily read and understand. c) Consistency. The comments should follow a standard style/format for better code reading. Code summarization is a task that tries to comprehend code and automatically generate descriptions directly from the source code. The summarization of code can also be viewed as a form of document expansion. Successful code summarization can

not only benefit the maintenance of source codes [4], [5], but also be used to improve the performance of code search using natural language queries [6], [7] and code categorization [8].

Motivation. Recent research has made inroads towards automatic generation of natural language descriptions of software. Currently, most of existing code summarization methods learn the semantic representation of source codes based on statistical language models [9], [4], and then generate comments based on templates or rules [10]. With the development of deep learning, some neural translation models [11], [5], [12] have also been introduced for code summarization, which mainly follow an encoder-decoder framework. They generally employ recurrent neural networks (e.g., LSTM) to encode the code snippets and utilize another recurrent neural network to decode that hidden state to coherent sentences. These models are typically trained to maximize the likelihood of the next word on the assumption that previous words and ground-truth are given. These models are limited from two aspects: a) the code sequential and structural information is not fully utilized on feature representation, which is critical for code understanding. For example, given two simple expressions “ $f=a+b$ ” and “ $f=c+d$ ”, although they are quite different as two lexical sequences, they share the same structure. In programming analysis, the structure of code is always represented by abstractive syntax trees, as shown in Figure 1a. b) these models, also termed “teacher-forcing”, suffer from the exposure bias since in test time the ground-truth is missing and previously generated words from the trained model distribution are used to predict the next word. Figure 1 presents a simple illustration of the discrepancy among training and testing process in these classical encoder-decoder models. In the testing phase, this exposure bias makes error accumulated and makes these models sub-optimal, not able to generate those words which are appropriate but with low probability to be drawn in the training phase.

Contribution. In this paper, we aim to address these two mentioned issues. To effectively capture the syntactic (or structure) information of code snippets, we employ abstract syntax tree (AST) [13], a data structure widely used in compilers, to represent the structure of program code. Figure 1a shows an example of Python code snippet and its corresponding AST. The root node is a composite node of type `FunctionDef`, leaf nodes which are typed as `Name` are tokens of code snippets. It’s worth mentioning that the tokens from AST parsing may be different from that from word segmentation. In our paper, we consider both of them. We parse the code snippets into abstract syntax trees (ASTs), and then propose an AST-based

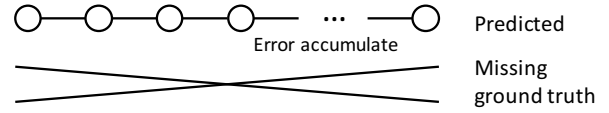

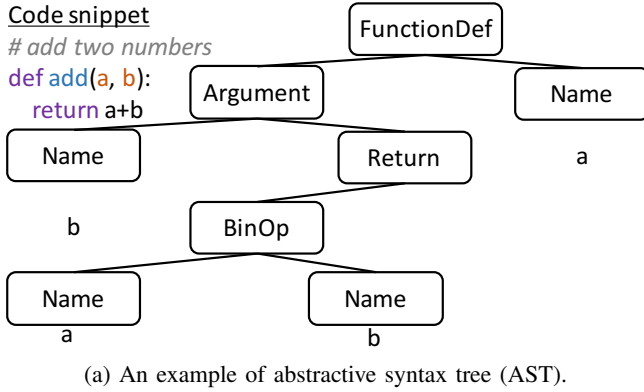


Figure 1: An illustration of the motivation of our paper. Traditional methods suffer from the following two limitations: a) on representing the code, the structure information of code is always ignored. b) traditional maximum likelihood based methods suffer from the exposure bias issue.

LSTM model [18] to represent the structure of code. We also use another LSTM model [42] to represent the sequential information of code. Besides, we apply a hybrid attention layer to fuse the structure representation and sequential representation of code on predicting the word, considering the alignment between predicting word and source word.

To overcome the exposure bias, we draw on the insights of deep reinforcement learning, which integrates exploration and exploitation into a whole framework. Instead of learning a sequential recurrent model to greedily look for the next correct word, we utilize an actor network and a critic network to jointly determine the next best word at each time step. The actor network, which provides the confidence of predicting the next word according to current state, serves as a local guidance. The critic network, which evaluates the reward value of all possible extensions of the current state, serves as a global and lookahead guidance. Our framework is able to include the good words that are with low probability to be drawn by using the actor network alone. To learn these two networks more efficiently, we start with pretraining an actor network using standard supervised learning with cross entropy loss, and pretraining a critic network with mean square loss. Then, we update the actor and critic networks according to the advantage reward composed of BLEU metric via policy gradient. We summarize our main contributions as follows.

- We propose a more comprehensive representation method for source code, with one AST-based LSTM for the structure of source code, and another LSTM for the sequential content of source code. Furthermore, a hybrid attention layer is applied to fuse these two representations.
- We propose an actor-critic network learning framework, an advanced deep reinforcement learning framework, to cope with the exposure bias issue existing in most traditional maximum likelihood estimation-based text generation framework.
- We validate our proposed model on a real-world dataset of 108,726 Python code snippets. Comprehensive experi-

ments show the effectiveness of the proposed model when compared with some state-of-the-art ones. To facilitate other researchers to repeat our experiments, we will release our dataset and source code later.

Organization. The remainder of this paper is organized as follows. We provide some background knowledge on neural language model, RNN encoder-decoder model and reinforcement learning in Section II for a better understanding of our proposed model. We also give an overview our proposed framework in Section II. Section III presents a hybrid embedding approach for code representation. Section IV shows our proposed deep reinforcement learning framework. Section V describes the dataset used in our experiment and shows the experimental results and analysis. Section VII highlights some works related to this paper. Section VI shows some threats to validity and limitations existing in our model. Finally, we conclude this paper and give some future research directions in Section VIII.

II. BACKGROUND

In this section, we first present some background knowledge on text generation used in this paper. To start with, we introduce some basic notations and terminologies. We consider the problem of learning to produce an output sequence $\mathbf{y} = (y_1, \dots, y_{|\mathbf{y}|})$, given an input $\mathbf{x} = (x_1, x_2, \dots, x_{|\mathbf{x}|})$, where $|\cdot|$ denotes the length of sequence. We will often use notation $\mathbf{y}_{f\dots l}$ to refer to subsequences of the form (y_f, \dots, y_l) . Two sets of input-output pairs (\mathbf{x}, \mathbf{y}) are assumed to be available for both training and testing.

A. Language Model

Language models compute the probability of occurrence of a number of words in a particular sequence. The probability of a sequence of T words $\{y_1, \dots, y_T\}$ is denoted as $p(y_1, \dots, y_T)$. Since the number of words coming before a word, y_i , varies depending on its location in the input document, $P(y_1, \dots, y_T)$

is usually conditioned on a window of n previous words rather than all previous words:

$$p(y_1, \dots, y_T) = \prod_{i=1}^{i=T} p(y_i | y_{1:i-1}) \approx \prod_{i=1}^{i=T} p(y_i | y_{i-(n-1):i-1}) \quad (1)$$

This kind of n-grams approach suffers apparent limitations [14], [15]. For example, the n-gram model probabilities are not derived directly from the frequency counts, because models derived this way have severe problems when confronted with any n-grams that have not explicitly been seen before. Furthermore, since they are simply smoothed counts of term co-occurrences, they are limited in their ability to generalize beyond the explicit features observed in training [16], [17].

The neural language model is a language model based on neural networks. Unlike the n-gram model which predicts a word based on a fixed number of predecing words, a neural language model can predict a word by predecing words with longer distances. Figure 2(a) shows the basic structure of an RNN. The neural network includes three layers, that is, an input layer which maps each word to a vector, a recurrent hidden layer which recurrently computes and updates a hidden state after reading each word, and an output layer which estimates the probabilities of the following word given the current hidden state. The RNN reads the words in the sentence one by one, and predicts the possible following word at each time step. At step t , it estimates the probability of the following word $p(y_{t+1} | y_{1:t})$ by three steps: First, the current word y_t is mapped to a vector e_t by the input layer: $e_t = \text{input}(y_t)$. Then, it generates the hidden state (values in the hidden layer) h_t at time t according to the previous hidden state h_{t-1} and the current input y_t :

$$h_t = f(h_{t-1}, y_t) \quad (2)$$

Finally, the $p(y_{t+1} | y_{1:t})$ is predicted according to the current hidden state h_t :

$$p(y_{t+1} | y_{1:t}) = g(h_t) \quad (3)$$

where g is a stochastic output layer (typically a softmax for discrete outputs) that generates output token.

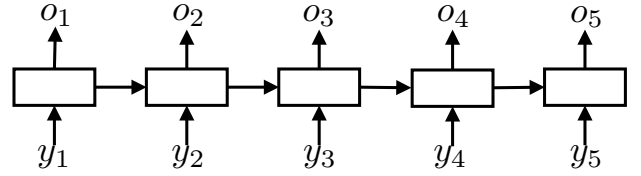
B. Attentional RNN Encoder-Decoder Model

RNN Encoder-Decoder has two recurrent neural networks (RNNs). The encoder transforms the code snippet \mathbf{x} into a sequence of hidden states $(\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_{|\mathbf{x}|})$ with a Recurrent Neural Network (RNN), while the decoder uses another RNN to generate one word y_{t+1} at a time in the target space. Generation is conditioned on all previously generated words $y_{1:t}$ and a dynamically created context vector \mathbf{c}_t , which encodes the source sentence:

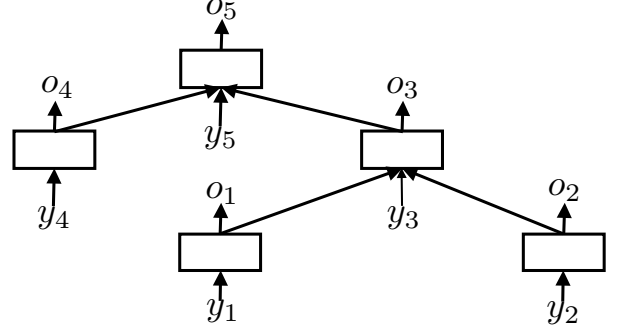
1) *Encoder*: As an RNN, the encoder has a hidden state \mathbf{h} , which is a fixed-length vector. At a time step t , the encoder computes the hidden state \mathbf{h}_t by:

$$h_t = f(\mathbf{h}_{t-1}, \mathbf{c}_{t-1}, e(y_t)). \quad (4)$$

Two common options for f are long short-term memory (LSTM) [42] and the gated recurrent unit (GRU) [35]. The



(a) The structure of recurrent neural network (RNN)



(b) The structure of tree-structured recurrent neural network (Tree-RNN)

Figure 2: LSTM and Tree-LSTM (adapted from [18]).

last symbol of y should be an end-of-sequence ($jeos_i$) symbol which notifies the encoder to stop and output the final hidden state \mathbf{h}_T , which is used as a vector representation of y .

2) *Decoder*: The output of the decoder is the target sequence $\mathbf{y} = (y_1, \dots, y_T)$. One input of the decoder is a $start_i$ symbol denoting the beginning of the target sequence. At a time step t , the decoder computes the hidden state \mathbf{h} and the conditional distribution t of the next symbol y_t by:

$$y_t \sim g(\mathbf{h}_{t-1}, \mathbf{c}_{t-1}) \quad (5)$$

where \mathbf{c}_t is the distinct context vector for y_t , and can be computed by:

$$\mathbf{c}_t = \sum_{j=1}^L \alpha_{t,j} \mathbf{h}_j \quad (6)$$

3) *Training goal*: The encoder and the decoder are jointly trained to maximize the conditional log-likelihood:

$$\max_{\theta} \mathcal{L}(\theta) = \max_{\theta} \frac{1}{N} \sum_{i=1}^N \log p(y_i | x_i; \theta), \quad (7)$$

where θ is the set of the model parameters; N is the size of the training set; and each (x_i, y_i) is a pair of a source sequence and a target sequence in the training set.

C. Reinforcement Learning for Better Decoding

We can see that this classical encoder-decoder framework targets on maximizing the likelihood of ground-truth word conditioned on previously generated words. Based on this framework, we propose our reinforcement learning framework since the text generation process can be viewed as a Markov Decision Process (MDP) $\{S, A, P, R, \gamma\}$. In the MDP setting, state \mathbf{s}_t at time step t consists of the source code snippets \mathbf{x} and the words/actions predicted until t , y_0, y_1, \dots, y_t . The

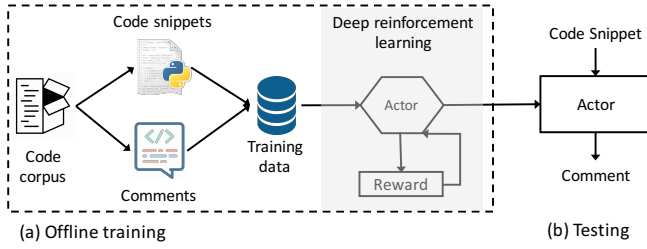


Figure 3: An overall workflow our proposed method.

action space is the dictionary \mathcal{Y} that the words are drawn from, *i.e.*, $y_t \subset \mathcal{Y}$. With the definition of the state, the state transition function P is $\mathbf{s}_{t+1} = \{s_t, y_{t+1}\}$, where the action y_{t+1} becomes a part of the next state \mathbf{s}_{t+1} and the reward r_{t+1} is received. $\gamma \in [0, 1]$ is the discount factor. The objective of generation process is to find a policy that maximizes the expected reward of generation sentence sampled from the model’s policy:

$$\max_{\theta} \mathcal{L}(\theta) = \max_{\theta} \mathbb{E}_{\substack{x \sim \mathcal{D} \\ \hat{y} \sim P_{\theta}(\cdot|x)}} [R(\hat{y}, x)], \quad (8)$$

where θ is the parameter of policy needed to be learnt, \mathcal{D} is the training set, \hat{y} is the predicted action/word, and R is the reward function. Our problem can be formulated as follows.

Given a code snippet $\mathbf{x} = (x_1, x_2, \dots, x_{|\mathbf{x}|})$, our goal is to find a policy that generates a sequence of words $\mathbf{y} = (y_1, y_2, \dots, y_{|\mathbf{y}|})$ from dictionary \mathcal{Y} with the objective of maximizing the expected reward.

D. An Overview of Proposed Framework

In this paper, we better represent the code, we propose a unified encoder-decoder framework with deep reinforcement learning for better decoding. Figure 3 shows the overall workflow of proposed method. It includes an offline training stage and an online summarization stage. In the training stage, we prepare a large-scale corpus of annotated $\langle \text{code}, \text{comment} \rangle$ pairs. The annotated pairs are used to train deep reinforcement learning model, *i.e.*, the actor-critic network as described in Section IV. After finished training, given a code snippet, corresponding comment can be generated trained actor network. The details of our method are explained in the following sections.

III. A HYBRID EMBEDDING APPROACH FOR CODE REPRESENTATION

As shown in Figure 4, our model contains an actor network and a critic network with sharing an encoder. In this section, we present a hybrid embedding approach for code representation. We apply a LSTM to represent the lexical level of code, and a Tree-LSTM to represent the syntactic level of code. We will elaborate our proposed actor-critic network in the next section.

A. Lexical Level

The key insight into lexical level representation of source code is that the comments are always extracted from the lexical of code, mainly from the function name, variable name and so on. It’s apparent that we adopt a LSTM to represent the lexical information of source code. # TODO: add some sentences.

B. Syntactic Level

Different from previous methods that just utilize sequence words to represent code, we also consider the structure information of source code. The front end of a compiler decomposes a program into constituents and produces intermediate code according to the syntax of the language [19]. These constituents are called programming constructs, and a context-free grammar specifies the syntax of programming constructs [19]. The AST is one type of intermediate code that represents the hierarchical syntactic structure of a program [19]. Ultimately, our goal is to specify learning-based techniques for encoding arbitrarily long sequences of lexical elements. Since the non-terminal nodes in ASTs subsume sequences of lexical elements [19], suppose each AST node has a special attribute *repr* that stores a vector representation, a code that characterizes the node and, by extension, the sequence of lexical elements the node subsumes. We mine the codes in such a way that similar sequences have similar codes. One learning-based technique is based on the AST, a tree representation that can have an arbitrary number of levels comprising nodes with an arbitrary number of children, but herein lies the problem.

Similar to a traditional LSTM unit, we propose AST-based LSTM where the LSTM unit also contains an input gate, a memory cell and an output gate. However, different from a standard LSTM unit which only has one forget gate for its previous unit, an AST-based LSTM unit contains multiple forget gates. Given an AST, for any node j , let the hidden state and memory cell of its l th child be h_{jl} and c_{jl} respectively. Refer to [18], the hidden state is updated as follows.

$$\begin{aligned} \mathbf{i}_j &= \sigma(\mathbf{W}^{(i)} \mathbf{x}_j + \sum_{l=1}^N \mathbf{U}_l^{(i)} \mathbf{h}_{jl} + \mathbf{b}^{(i)}), \\ \mathbf{f}_{jk} &= \sigma(\mathbf{W}^{(f)} \mathbf{x}_j + \sum_{l=1}^N \mathbf{U}_{kl}^{(f)} \mathbf{h}_{jl} + \mathbf{b}^{(f)}), \\ \mathbf{o}_j &= \sigma(\mathbf{W}^{(o)} \mathbf{x}_j + \sum_{l=1}^N \mathbf{U}_l^{(o)} \mathbf{h}_{jl} + \mathbf{b}^{(o)}), \\ \mathbf{u}_j &= \tanh(\mathbf{W}^{(u)} \mathbf{x}_j + \sum_{l=1}^N \mathbf{U}_l^{(u)} \mathbf{h}_{jl} + \mathbf{b}^{(u)}), \\ \mathbf{c}_j &= \mathbf{i}_j \odot \mathbf{u}_j + \sum_{l=1}^N \mathbf{f}_{jl} \odot \mathbf{c}_{jl}, \\ \mathbf{h}_j &= \mathbf{o}_j \odot \tanh(\mathbf{c}_j) \end{aligned} \quad (9)$$

where $k = 1, 2, \dots, N$. Each of \mathbf{i}_j , \mathbf{f}_{jk} , \mathbf{o}_j and \mathbf{u}_j denotes an input gate, a forget gate, an output gate, and a state for updating the memory cell, respectively. $\mathbf{W}^{(\cdot)}$ and $\mathbf{U}^{(\cdot)}$ are

weight matrices, $\mathbf{b}^{(\cdot)}$ is a bias vector, and \mathbf{x}_j is the word embedding of the j -th node. $\sigma(\cdot)$ is the logistic function, and the operator \odot denotes element-wise multiplication between vectors. It's worth mentioning that when the tree is simply a chain, namely $N = 1$, the AST-based LSTM unit reduces to the standard LSTM.

Notice that the number of children N varies for different nodes of different ASTs, which may cause problem in parameter-sharing. For simplification, we transform the generated ASTs to binary trees by the following two steps which have been adopted in [20]: a) split nodes with more than 2 children, generate a new right child together with the old left child as its children, and then put all children except the leftmost as the children of this new node. Repeat this operation in a top-down way until only nodes with 0, 1, 2 children left; b) combine nodes with 1 child with its child.

IV. DEEP REINFORCEMENT LEARNING FOR CODE SUMMARIZATION

In this section, we adopt the advanced deep learning framework named actor-critic network, which has been successfully used in the AlphaGo [47]. We introduce the actor and critic network respectively and then present how to train the two networks simultaneously.

A. Actor Network

Different parts of the code make different contribution to the final output of comments. We adopt an attention mechanism [21] which has been successfully used in neural machine translation. In the attention layer, we have two attention scores, one $\alpha_t^{str}(j)$ for structural representation and another $\alpha_t^{txt}(j)$ for sequential representation of code. At t -th step of the decoder process, the attention scores $\alpha_t^{str}(j)$ and $\alpha_t^{txt}(j)$ are calculated as follows:

$$\alpha_t^{str}(j) = \frac{\exp(\mathbf{h}_j^{str} \cdot \mathbf{s}_t)}{\sum_{k=1}^n \exp(\mathbf{h}_k^{str} \cdot \mathbf{s}_t)}, \quad (10)$$

$$\alpha_t^{txt}(j) = \frac{\exp(\mathbf{h}_j^{txt} \cdot \mathbf{s}_t)}{\sum_{k=1}^n \exp(\mathbf{h}_k^{txt} \cdot \mathbf{s}_t)}, \quad (11)$$

where n is the number of code tokens; $\mathbf{h}_j \cdot \mathbf{s}_t$ is the inner product of \mathbf{h}_j and \mathbf{s}_t , which is used to directly calculate the similarity score between \mathbf{h}_j and \mathbf{s}_t . The t -th context vector \mathbf{d}_t is calculated as the summarization vector weighted by $\alpha_t(j)$:

$$\mathbf{d}_t^{str} = \sum_{j=1}^n \alpha_t^{str}(j) \mathbf{h}_j^{str}, \quad (12)$$

$$\mathbf{d}_t^{txt} = \sum_{j=1}^n \alpha_t^{txt}(j) \mathbf{h}_j^{txt}. \quad (13)$$

To integrate the structural context vector and the textual vector, we concatenate them firstly and then feed them into a one-layer linear network:

$$\mathbf{d}_t = \mathbf{W}_{\mathbf{d}_t} [\mathbf{d}_t^{str}; \mathbf{d}_t^{txt}] + \mathbf{b}_{\mathbf{d}_t}, \quad (14)$$

where $[\mathbf{d}_t^{str}; \mathbf{d}_t^{txt}]$ is the concatenation of \mathbf{d}_t^{str} and \mathbf{d}_t^{txt} . The context vector is then used for the $(t+1)$ -th word prediction by putting an additional hidden layer $\tilde{\mathbf{s}}_t$:

$$\tilde{\mathbf{s}}_t = \tanh(\mathbf{W}_c [\mathbf{s}_t; \mathbf{d}_t] + \mathbf{b}_d), \quad (15)$$

where $[\mathbf{s}_t; \mathbf{d}_t]$ is the concatenation of \mathbf{s}_t and \mathbf{d}_t . The model predicts the t -th word by using a softmax function. Let p_π denote a policy π determined by the actor network, $p_\pi(y_t | \mathbf{s}_t)$ denote the probability distribution of generating t -th word y_t , we can get the following equation:

$$p_\pi(y_t | \mathbf{s}_t) = \text{softmax}(\mathbf{W}_s \tilde{\mathbf{s}}_t + \mathbf{b}_s), \quad (16)$$

B. Critic Network

Unlike traditional encoder-decoder framework which generates sequence directly via maximizing the likelihood of next word given the ground truth word, we directly optimize the evaluation metrics such as BLEU [22] for code summarization. We apply a critic network to approximate the value of generated actions at time step t . Different from the actor network, this critic network outputs a single value instead of a probability distribution on each decoding step. Before we introduce our value network, we firstly introduce the value function.

Given the policy π , sampled actions and reward function, the value function V^π is defined as the prediction of total reward from the state \mathbf{s}_t at step t under policy π , which is formulated as follows:

$$V^\pi(\mathbf{s}_t) = \mathbb{E}_{\mathbf{s}_{t+1:T}, y_{t:T}} \left[\sum_{l=0}^{T-t} r_{t+l} | y_{t+1}, \dots, y_T, \mathbf{h} \right], \quad (17)$$

where T is the max step of decoding; \mathbf{h} is the representation of code snippet. For code summarization, we can only obtain an evaluation score (BLEU) when the sequence generation process (or episode) is finished. The episode terminates when step exceeds the max-step T or generating the end-of-sequence (EOS) token. Therefore, we define the reward as follows:

$$r_t = \begin{cases} 0 & t < T \\ BLEU & t = T \text{ or } EOS \end{cases}. \quad (18)$$

Mathematically, the critic network tries to minimize the following loss function, when mean square error is used.

$$\mathcal{L}(\phi) = \frac{1}{2} \|V^\pi(\mathbf{s}_t) - V_\phi^\pi(\mathbf{s}_t)\|^2, \quad (19)$$

where $V^\pi(\mathbf{s}_t)$ is the target value, $V_\phi^\pi(\mathbf{s}_t)$ is the value predicted by critic network and ϕ is the parameter of critic network.

C. Model Training

We use the policy gradient method to optimize policy directly, which is widely used in reinforcement learning. For actor network, the goal of training is to minimize the negative expected reward, which can be defined as $\mathcal{L}(\theta) = -\mathbb{E}_{y_1, \dots, y_T \sim \pi} (\sum_{l=t}^T r_l)$, where θ is the parameter of actor network. Denote all the

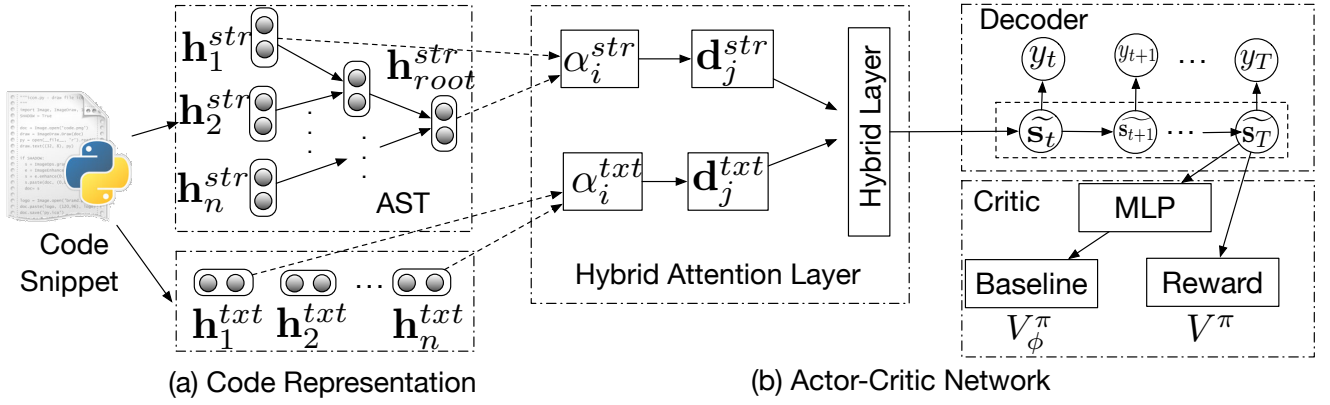


Figure 4: An overview of our proposed tree-structured actor-critic network learning framework for code summarization. The input code snippet is parsed into abstract syntax tree and sequential tokens firstly. The actor and critic network share a code encoder network. (a) The actor network computes the probability of executing an action y_t at a certain state s_t , by $p_\pi(y_t|s_t)$. (b) Given a state s_t , the value network evaluates its value (baseline), the advantage is defined as $|baseline - reward|$.

parameters as $\Theta = \{\theta, \phi\}$, the total loss of our model can be represented as $\mathcal{L}(\Theta) = \mathcal{L}(\theta) + \mathcal{L}(\phi)$.

For policy gradient, it is typically better to train an expression of the following form according to [23]:

$$\nabla_\theta \mathcal{L}(\Theta) = \mathbb{E} \left[\sum_{t=0}^{T-1} A^\pi(s_t, y_{t+1}) \nabla_\theta \log \pi_\theta(y_{t+1}|s_t) \right], \quad (20)$$

where $A^\pi(s_t, y_{t+1})$ is advantage function. The reason why we choose advantage function is that it achieves smaller variance when compared with some other ones such as TD residual and reward with baseline.

According to the definition of advantage function, we can formulate the advantage function as follows. One can refer to [24] for more details.

$$A^\pi(s_t, y_t) = Q^\pi(s_t, y_t) - V^\pi(s_t), \quad (21)$$

where $Q^\pi(s_t, y_t)$ is the state-action value function defined as $Q^\pi(s_t, y_t) = \mathbb{E}_{s_{t+1:T}, r_{t+1:T}} \left[\sum_{l=0}^{T-t} r_{t+l} \right]$. From this formulation, we can find that the advantage function measures whether or not the action is better or worse than the policy's default behavior. Therefore, a step in the policy gradient direction can increase the probability of better-than-average actions and decrease the probability of worse-than-average actions.

On the other hand, the gradient of critic network is calculated as follows:

$$\nabla_\phi \mathcal{L}(\Theta) = \sum_{t=0}^{T-1} [V^\pi(s_t) - V_\phi^\pi(s_t)] \nabla_\phi V_\phi^\pi(s_t). \quad (22)$$

We employ stochastic gradient descent with the diagonal variant of AdaGrad [25] to optimize the parameters of our framework. At the t -th step, the parameters Θ are updated by:

$$\Theta \leftarrow \Theta - \frac{\rho}{\sqrt{\sum_{i=1}^t g_i^2}} g_t \quad (23)$$

Algorithm 1 Actor-Critic training for code summarization.

- 1: Initialize actor $\pi_{y_{t+1}|s_t}$ and critic $V^\pi(s_t)$ with random weights θ and ϕ ;
 - 2: Pre-train the actor to predict ground truth y_t given $\{y_1, \dots, y_{t-1}\}$ by minimizing Eq. ??;
 - 3: Pre-train the critic to estimate $V(s_t)$ by with fixed actor;
 - 4: **for** $t = 1 \rightarrow T$ **do**
 - 5: Receive a random example, and generate sequence of actions $\{y_1, \dots, y_T\}$ according to current policy π_θ ;
 - 6: Calculate advantage estimate A^π according to Eq. 21;
 - 7: Update critic weights ϕ using the gradient in Eq. 22;
 - 8: Update actor weights θ using the gradient in Eq. 20.
-

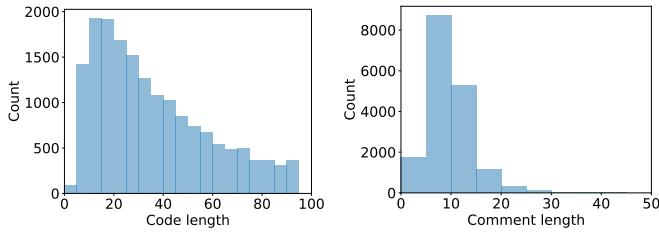
where ρ is the initial learning rate and g_t is the subgradient at time t . Algorithm 1 summarizes our proposed model described above.

V. EXPERIMENTS AND ANALYSIS

To evaluate our proposed approach, in this section, we conduct experiments to answer the following questions:

- **RQ1.** Does our proposed approach improve the performance of code summarization when compared with some state-of-the-art approaches.
- **RQ2.** What's the effectiveness of each component for our proposed model? For example, what about the performance of hybrid code representation and the that of reinforcement learning?
- **RQ3.** What's the performance of our proposed model on the datasets with different code or comment length?

We ask RQ1 to evaluate our deep reinforcement learning-based model compared to some state-of-the-art baselines, which we describe in the following subsection. We ask RQ2 in order to evaluate each component of our model. We ask RQ3 to evaluate our model when varying the length of code or comment. In the following subsections, we first describe the dataset, some



(a) Code length distribution. (b) Comment length distribution.

Figure 5: Length distribution of testing data.

evaluation metrics and the training details. Then, we introduce the baseline for RQ1. Finally, we report our results and analysis for the research questions.

A. Dataset Preparation

We evaluate the performance of method using the dataset in [26], which is obtained from a popular open source projects hosting platform, GitHub. The dataset contains 108,726 code-comment pairs. The vocabulary size of code and comment is 50,400 and 31,350, respectively. For cross-validation, we shuffle the dataset firstly and use the first 60% for training, 20% for validation and the left for testing. To construct the tree-structure of code, we parse Python code into abstract syntax trees via `ast`¹ lib. To convert code into sequential text, we tokenize the code by `{ ' " ' ; :) (! ? (space) }`, which has been used in [8]. We tokenize the comment by `{(space)}`.

Figure 5 shows the length distribution of code and comment, respectively. From Figure 5a, we can find that the length of most code snippets are located between 20 to 60. This verify the quote in [27] that “Functions should hardly ever be 20 lines long”. In Python language, the limited length should be shorter. From Figure 5b, we can notice that the length of nearly all the comments are between 5 to 15. This reveals the comment sequence that we need to generate will not be too long.

B. Evaluation metrics.

We evaluate the performance of our proposed model based on four widely-used evaluation criteria in the area of neural machine translation and image captioning, *i.e.*, BLEU [22] METEOR [28], ROUGE-L[29] and CIDER[30]. BLEU measures the average n-gram precision on a set of reference sentences, with a penalty for overly short sentences. METEOR is recall-oriented and measures how well our model captures content from the references in our output. ROUGE-L takes into account sentence level structure similarity naturally and identifies longest co-occurring in sequence n-grams automatically. CIDER is a consensus based evaluation protocol for image descriptions. To make the paper be compacted, we put the formulation of each metric in Table IV (see Appendix A).

C. Training Details

The hidden size of the encoder and decoder LSTM networks are both set to be 512, and the word embedding size is set to be

¹<https://docs.python.org/2/library/ast.html>

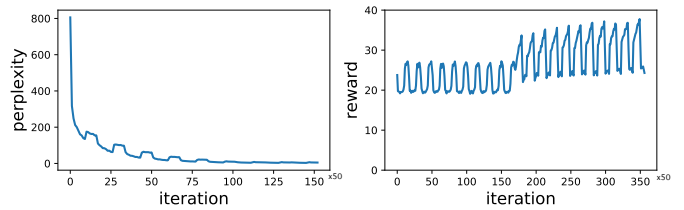


Figure 6: Iteration of training perplexity and reward.

512. The mini-batch size is set to be 64, while the learning rate is set to be 0.001. We pretrain both actor network and critic network with 10 epochs each, and train the actor-critic network simultaneously 10 epoches. We record the perplexity²/reward every 50 iterations. Figure 6 shows the perplexity and reward curves of our method.

All the experiments in this paper are implemented with Python 2.7, and run on a computer with an 2.2 GHz Intel Core i7 CPU, 64 GB 1600 MHz DDR3 RAM, and a Titan X GPU with 12 GB memory, running Ubuntu 16.04.

D. RQ1: Compared to Baselines

We compare our model with the following baseline methods:

- Seq2Seq [31] is a classical encoder-decoder framework in neural machine translation, which encodes the source sentences into a hidden space, and decodes it into target ones. In our comparison, the encoder and decoder are both based on LSTM.
- Seq2Seq+Attn [21] is a derived version of Seq2Seq model with an attentional layer for sentence alignment.
- Tree2Seq [20] follows the same architecture with Seq2Seq and applies Tree-LSTM as encoder for the task of code clone detection.
- Tree2Seq+Attn [32] is a derived version of Tree2Seq model with an attentional layer, which has been applied in neural machine translation
- Hybrid2Seq(+Attn+A2C) represents three versions of our proposed model with/without Attn/A2C component.

Table I shows the experimental results of comparison between our proposed model and some previous ones. From this table, we can find that our proposed model outperforms other baselines in almost all of evaluation metrics. When comparing Seq2Seq/Tree2Seq with its correspond attention-based version, we can see that attention is really effective in aligning the code tokens with comment tokens. We can also find that the performance of simply encoding the tree structure of code is worse than that of simply encoding the code as sequence. This can be illustrated by that the words of comments are always drawn from the tokens of code. Thus, our model which considers both the structure and sequential information of code achieves the best performance in this comparison.

²Perplexity is a function of cross entropy loss, which has been widely used in evaluation of many natural language processing tasks.

Table I: Comparison of the overall performance between our model and previous methods. (Best scores are in boldface.)

	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE-L	CIDER
Seq2Seq	0.1660	0.0251	0.0100	0.0056	0.0535	0.2838	0.1262
Seq2Seq+Attn	0.1897	0.0419	0.0200	0.0133	0.0649	0.3083	0.2594
Tree2Seq	0.1649	0.0236	0.0096	0.0053	0.0501	0.2794	0.1168
Tree2Seq+Attn	0.1887	0.0417	0.0197	0.0129	0.0644	0.3068	0.2331
Hybrid2Seq+Attn+A2C (Our)	0.2527	0.1033	0.0640	0.0441	0.0929	0.3913	0.7501

Table II: Effectiveness of each component for our proposed model. (Best scores are in boldface.)

	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE-L	CIDER
Seq2Seq+Attn+A2C	0.2421	0.0919	0.0513	0.0325	0.0882	0.3935	0.6390
Tree2Seq+Attn+A2C	0.2309	0.0854	0.0499	0.0338	0.0843	0.3767	0.6060
Hybrid2Seq	0.1837	0.0379	0.0183	0.0122	0.0604	0.3020	0.2223
Hybrid2Seq+Attn	0.1965	0.0516	0.0280	0.0189	0.0693	0.3154	0.3475
Hybrid2Seq+Attn+A2C (Our)	0.2527	0.1033	0.0640	0.0441	0.0929	0.3913	0.7501

E. RQ2: Performance of Hybrid Code Representation and Reinforcement Learning

Table II shows the effectiveness of some main components in our proposed model. From this table, comparing the results of Seq2Seq+Attn/Tree2Seq+Attn with and without (Table I) the advantage actor-critic (A2C) network, we can see that the proposed A2C component can really boost the performance of comment generation for source code. We can also find the proposed approach of integrating the LSTM for content and AST-based LSTM for structure is effective on representing the code as compared with the corresponding non-hybrid ones in Table I. Furthermore, it also verifies that our proposed hybrid attention mechanism works well in our model.

F. RQ4: Parameter Analysis

Figure 7 and Figure 8 show the performance of our proposed method when compared with two baselines on datasets of varying code length and comment lengths, respectively. From Figure 7, we can see that our model performs best when compared with other baselines on four metrics with respect to different code length. Additionally, we can see that the our proposed model has a stable performance even though the code length increase dramatically. We attribute this effect to the hybrid representation we adopt in our model.

From Figure 8, remind that comment length distribution in Figure 5b. Since nearly all the comment length of testing data are under 20, we ignore the performance analysis over the data whose comment length are larger than 20. From this figure, we can see the performances our model and baselines vary dramatically on four metrics with respect to different comment length.

G. Qualitative Analysis and Visualization

We show two examples in Table III. It’s clear that the generated comments by our model is closest to the ground truth. Although those models without A2C can generate some tokens which are also in the ground truth, they can’t predict those tokens which are not frequently appeared in the training data. On the contrary, our actor-critic learning based model can generate some tokens which are more closer to the ground truth, like `git`, `symbolic`. This can be illustrated by the fact

that our model has a more comprehensive exploration on the word space and optimizes the BLEU score directly.

In Table III, we also visualize two attentions in our proposed model for the target sentences. For example, for Case 1 with target sentence `check if git is installed .`, we can notice that the str-attn (left of figure) focuses more on tokens like `0SError`, `False`, `git`, `version`, which represent the structure of code. On the other hand, the attention of txt-attn (right of figure) is comparatively dispersed, and have a focus on some tokens like `def`, which is of little significance for code summarization. This verifies our assumption that LSTM can capture the sequential content of code, and AST-based LSTM can capture the structure information of code. Thus, it’s reasonable to fuse them together for a better representation.

VI. THREATS TO VALIDITY AND LIMITATIONS

One threat to validity is that our approach is experimented only on Python code collected from GitHub, so they may not be representative of all the comments. However, Python is a popular programming language used in a large number of projects. In the future, we will extend our approach to other programming languages.

Another threat to validity is on the metrics we choose for evaluation. It has always been a tough challenge to evaluate the similarity between two sentences for the tasks such as neural machine translation, image captioning [1]. In this paper, we only adopt four popular metrics, it is necessary for us to evaluate the performance of generated text from more perspectives, such as human evaluation.

Furthermore, in the deep reinforcement learning perspective, we only set the BLEU score of generated sentence as the reward. Its well known that for a reinforcement learning framework, one of the biggest challenge is how to design a reward function to measure the value of action correctly, and it is still an open problem. In our future work, we plan to devise a reward function that can reflect the value of each action more correctly.

VII. RELATED WORK

In this section, we briefly review some related work from the aspects of deep learning on code analysis, source code summarization and deep reinforcement learning.

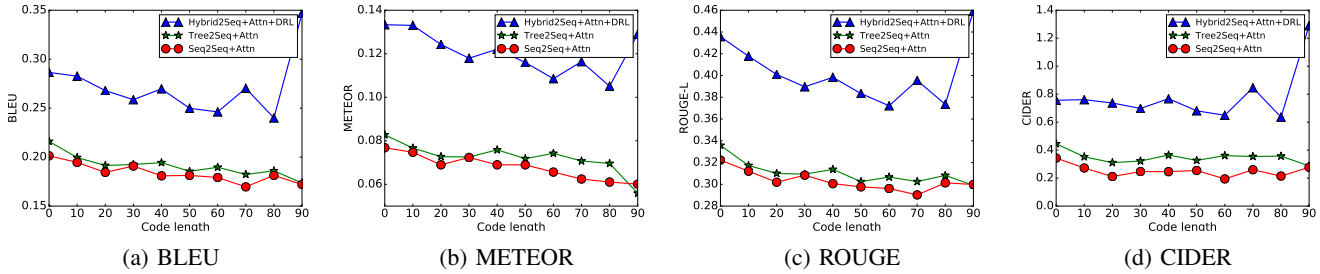


Figure 7: Experimental results of our proposed method and some baselines on different metrics with respect to varying code length.

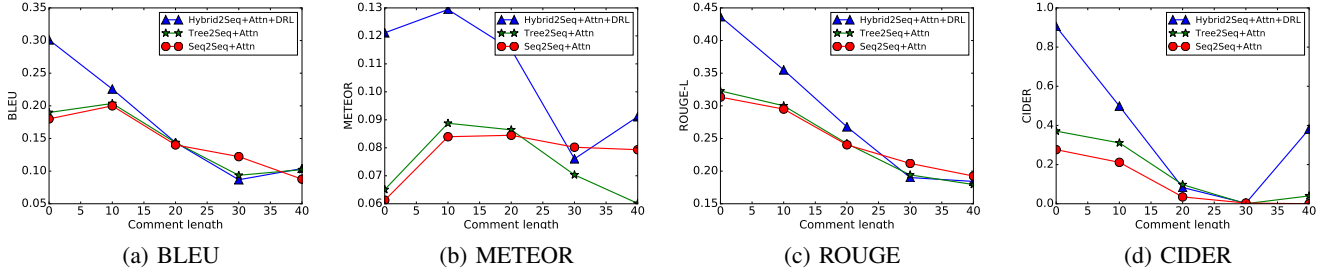


Figure 8: Experimental results of our proposed method and some baselines on different metrics with respect to varying comment length.

Table III: Examples of code summarization generated by each model.

	Case 1	Case 2
Code snippet	<pre>def _has_git(): try: subprocess.check_call([git, --version], stdout=subprocess.DEVNULL, stderr=subprocess.DEVNULL) except (OSError, subprocess .CalledProcessError): return False else: return True</pre>	<pre>def tensor3(name=None, dtype=None): if (dtype is None): dtype=config.floatX type=CudaNdarrayType(dtype=dtype, broadcastable= (False, False, False)) return type(name)</pre>
Ground truth	check if git is installed .	return a symbolic 3-d variable .
Seq2Seq	helper function to create a new figure manager instance .	yaml
Seq2Seq+Attn	return true if the user has access to the specified resource .	a decorator that returns a new class that will return a new class name .
Tree2Seq+Attn	test that validate_folders throws a foldermissingerror .	helper function for #4957 .
Hybrid2Seq+Attn	returns the number of git modules that are not installed .	return the path to the currently running server .
Hybrid2Seq+Attn+A2C	returns true if git is installed .	return a symbolic graph .
Attention visualization		

A. Deep Learning on Code Analysis

With the successful development of deep learning, it has also become more and more prevalent for representing source

code in the domain of software engineering research. Gu et al. [33] use a sequence-to-sequence deep neural network [31], originally introduced for SMT, to learn intermediate distributed vector representations of natural language queries

which they use to predict relevant API sequences. Mou et al. [34] learn distributed vector representations using custom convolutional neural networks to represent features of snippets of code, then they assume that student solutions to various coursework problems have been intermixed and seek to recover the solution-to-problem mapping via classification. Li et al. [35] learn distributed vector representations for the nodes of a memory heap and use the learned representations to synthesize candidate formal specifications for the code that produced the heap. Li et al. [35] exploit heap structure to define graph neural networks, a new machine learning model based on GRUs (a type of RNN) to directly learn from heap graphs. Piech et al. [36] and Parisotto et al. [37] learn distributed representations of source code input/output pairs and use them to assess and review student assignments or to guide program synthesis from examples.

Neural code-generative models of code also use distributed representations to capture context, a common practice in NLP. For example, the work of Maddison and Tarlow [38] and other neural language models (e.g. LSTMs in Dam et al. [39]) describe context distributed representations while sequentially generating code. Ling et al. [40] and Allamanis et al. [41] combine the code-context distributed representation with a distributed representations of other modalities (e.g. natural language) to synthesize code. While all of these representations can, in principle, encode unbounded context, handling all code dependencies of arbitrary length is an unsolved problem. Some neural architectures, such as LSTMs [42], GRUs [43] and their variants, have made progress on this problem and handle moderately long-range dependencies.

B. Source Code Summarization

Code summarization is a novel task in the area of software engineer, which has drawn great attention in recent years. The existing work for the problem of code summarization can be mainly categorized as rule based approaches [10], statistical language model based approaches [4] and deep learning based approaches [11], [5], [12]. Sridhara et al. [10] construct a software word usage model first, and generate comment according to the tokenized function/variable names via rules. Movshovitz-Attias et al. [4] predict comments from JAVA source files of open source projects using topic models and n-grams. In [11], the authors introduce an attentional neural network that employs convolution on the input tokens to detect local time-invariant and long-range topical attention features to summarize source code snippets into short, descriptive function name-like summaries. Iyer et al. [5] propose to use LSTM networks with attention to produce sentences that describe C# code snippets and SQL queries. In Haije's thesis [12], the code summarization problem is modeled as a translation task, and some classical translation models such as Seq2Seq [31] and Seq2Seq with attention [21] are employed. Unlike previous study, we take the tree structure and sequential content of source code into consideration for a better representation of code.

C. Deep Reinforcement Learning

Reinforcement learning [44], [45], [46], known as “a machine learning technique concerning how software agent ought to take actions in an environment so as to maximize some notion of cumulative reward”, is well suited for the task of decision-making. Recently, professional-level computer Go program has been designed by Silver et al. [47] using deep neural networks and Monte Carlo Tree Search. Human-level gaming control [48] has been achieved through deep Q-learning. A visual navigation system [49] has been proposed recently based on actor-critic reinforcement learning model.

Recently, with the success of AlphaGo[47], deep reinforcement learning has shown great potential for the task of decision-making through exploitation and exploration. Text generation can also be formulated as a decision-making problem and there have been several reinforcement learning-based works on this specific tasks, including image captioning [50], dialogue generation [51] and sentence simplification [52]. Ren et al. [50] propose an actor-critic deep reinforcement learning model with an embedding reward for image captioning. Li et al. [51] integrate a developer-defined reward with REINFORCE algorithm for dialogue generation. In this paper, we follow an actor-critic reinforcement learning framework, while our focus is on encoding the structural and sequential information of code snippets simultaneously with an attention mechanism.

VIII. CONCLUSION AND FUTURE WORK

In this paper, we propose a tree-structured actor-critic learning model to generate summarization for code snippets. Specifically, we first encode the structure and sequential content of code via AST-based LSTM and LSTM respectively. Then we add a hybrid attention layer to integrate them. We then feed the code representation vector into an actor-critic framework. Comprehensive experiments on a real-world dataset show that our proposed model outperforms other competitive baselines and achieves state-of-the-art performance on automatic metrics, namely BLEU, METEOR, ROUGE-L and CIDER.

In our future work, we plan to design a copy mechanism to cope with rare words which are out of our vocabulary, and extend our experiments to other programming languages such as Java.

REFERENCES

- [1] B. P. Lientz and E. B. Swanson, "Software maintenance management," 1980.
- [2] S. C. B. de Souza, N. Anquetil, and K. M. de Oliveira, "A study of the documentation essential to software maintenance," in *Proceedings of the 23rd annual international conference on Design of communication: documenting & designing for pervasive information*. ACM, 2005, pp. 68–75.
- [3] M. Kajko-Mattsson, "A survey of documentation practice within corrective maintenance," *Empirical Software Engineering*, vol. 10, no. 1, pp. 31–55, 2005.
- [4] D. Movshovitz-Attias and W. W. Cohen, "Natural language models for predicting programming comments," 2013.
- [5] S. Iyer, I. Konstas, A. Cheung, and L. Zettlemoyer, "Summarizing source code using a neural attention model," in *ACL (1)*, 2016.
- [6] D. Yang, A. Hussain, and C. V. Lopes, "From query to usable code: An analysis of stack overflow code snippets," in *Mining Software Repositories (MSR), 2016 IEEE/ACM 13th Working Conference on*. IEEE, 2016, pp. 391–401.
- [7] L. Nie, H. Jiang, Z. Ren, Z. Sun, and X. Li, "Query expansion based on crowd knowledge for code search," *IEEE Transactions on Services Computing*, vol. 9, no. 5, pp. 771–783, 2016.
- [8] A. T. Nguyen and T. N. Nguyen, "Automatic categorization with deep neural network for open-source java projects," in *Proceedings of the 39th International Conference on Software Engineering Companion*. IEEE Press, 2017, pp. 164–166.
- [9] Y. Oda, H. Fudaba, G. Neubig, H. Hata, S. Sakti, T. Toda, and S. Nakamura, "Learning to generate pseudo-code from source code using statistical machine translation (t)," in *Automated Software Engineering (ASE), 2015 30th IEEE/ACM International Conference on*. IEEE, 2015, pp. 574–584.
- [10] G. Sridhara, E. Hill, D. Muppaneni, L. Pollock, and K. Vijay-Shanker, "Towards automatically generating summary comments for java methods," in *Proceedings of the IEEE/ACM international conference on Automated software engineering*. ACM, 2010, pp. 43–52.
- [11] M. Allamanis, H. Peng, and C. Sutton, "A convolutional attention network for extreme summarization of source code," in *International Conference on Machine Learning*, 2016, pp. 2091–2100.
- [12] T. Haije, B. O. K. Intelligentie, E. Gavves, and H. Heuer, "Automatic comment generation using a neural translation model," 2016.
- [13] I. D. Baxter, A. Yahin, L. Moura, M. Sant'Anna, and L. Bier, "Clone detection using abstract syntax trees," in *Software Maintenance, 1998. Proceedings., International Conference on*. IEEE, 1998, pp. 368–377.
- [14] R. Rosenfeld, "Two decades of statistical language modeling: Where do we go from here?" *Proceedings of the IEEE*, vol. 88, no. 8, pp. 1270–1278, 2000.
- [15] A. Mnih and Y. W. Teh, "A fast and simple algorithm for training neural probabilistic language models," *arXiv preprint arXiv:1206.6426*, 2012.
- [16] Y. Bengio, R. Ducharme, P. Vincent, and C. Jauvin, "A neural probabilistic language model," *Journal of machine learning research*, vol. 3, no. Feb, pp. 1137–1155, 2003.
- [17] F. Morin and Y. Bengio, "Hierarchical probabilistic neural network language model," in *Aistats*, vol. 5. Citeseer, 2005, pp. 246–252.
- [18] K. S. Tai, R. Socher, and C. D. Manning, "Improved semantic representations from tree-structured long short-term memory networks," *arXiv preprint arXiv:1503.00075*, 2015.
- [19] A. V. Aho, R. Sethi, and J. D. Ullman, "Compilers, principles, techniques," *Addison Wesley*, vol. 7, no. 8, p. 9, 1986.
- [20] H. H. Wei and M. Li, "Supervised deep features for software functional clone detection by exploiting lexical and syntactical information in source code," 2017.
- [21] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.
- [22] K. Papineni, S. Roukos, T. Ward, and W. J. Zhu, "Bleu: a method for automatic evaluation of machine translation," in *Proceedings of the 40th annual meeting on association for computational linguistics*. Association for Computational Linguistics, 2002, pp. 311–318.
- [23] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, "High-dimensional continuous control using generalized advantage estimation," *arXiv preprint arXiv:1506.02438*, 2015.
- [24] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," in *International Conference on Machine Learning*, 2016, pp. 1928–1937.
- [25] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *Journal of Machine Learning Research*, vol. 12, no. Jul, pp. 2121–2159, 2011.
- [26] A. V. M. Barone and R. Sennrich, "A parallel corpus of python functions and documentation strings for automated code documentation and code generation," *arXiv preprint arXiv:1707.02275*, 2017.
- [27] R. C. Martin, *Clean code: a handbook of agile software craftsmanship*. Pearson Education, 2009.
- [28] S. Banerjee and A. Lavie, "Meteor: An automatic metric for mt evaluation with improved correlation with human judgments," in *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, vol. 29, 2005, pp. 65–72.
- [29] C. Y. Lin, "Rouge: A package for automatic evaluation of summaries," *Text Summarization Branches Out*, 2004.
- [30] R. Vedantam, C. Lawrence Zitnick, and D. Parikh, "Cider: Consensus-based image description evaluation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 4566–4575.
- [31] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Advances in neural information processing systems*, 2014, pp. 3104–3112.
- [32] A. Eriguchi, K. Hashimoto, and Y. Tsuruoka, "Tree-to-sequence attentional neural machine translation," *arXiv preprint arXiv:1603.06075*, 2016.
- [33] X. Gu, H. Zhang, D. Zhang, and S. Kim, "Deep api learning," in *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. ACM, 2016, pp. 631–642.
- [34] L. Mou, G. Li, L. Zhang, T. Wang, and Z. Jin, "Convolutional neural networks over tree structures for programming language processing," in *AAAI*, vol. 2, no. 3, 2016, p. 4.
- [35] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel, "Gated graph sequence neural networks," *arXiv preprint arXiv:1511.05493*, 2015.
- [36] C. Piech, J. Huang, A. Nguyen, M. Phulsuksombati, M. Sahami, and L. Guibas, "Learning program embeddings to propagate feedback on student code," *arXiv preprint arXiv:1505.05969*, 2015.
- [37] E. Parisotto, A.-r. Mohamed, R. Singh, L. Li, D. Zhou, and P. Kohli, "Neuro-symbolic program synthesis," *arXiv preprint arXiv:1611.01855*, 2016.
- [38] C. Maddison and D. Tarlow, "Structured generative models of natural source code," in *International Conference on Machine Learning*, 2014, pp. 649–657.
- [39] H. K. Dam, T. Tran, and T. Pham, "A deep language model for software code," *arXiv preprint arXiv:1608.02715*, 2016.
- [40] W. Ling, E. Grefenstette, K. M. Hermann, T. Kočíský, A. Senior, F. Wang, and P. Blunsom, "Latent predictor networks for code generation," *arXiv preprint arXiv:1603.06744*, 2016.
- [41] M. Allamanis, D. Tarlow, A. Gordon, and Y. Wei, "Bimodal modelling of source code and natural language," in *International Conference on Machine Learning*, 2015, pp. 2123–2132.
- [42] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [43] K. Cho, B. Van Merriënboer, D. Bahdanau, and Y. Bengio, "On the properties of neural machine translation: Encoder-decoder approaches," *arXiv preprint arXiv:1409.1259*, 2014.
- [44] R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning," in *Reinforcement Learning*. Springer, 1992, pp. 5–32.
- [45] V. R. Konda and J. N. Tsitsiklis, "Actor-critic algorithms," in *Advances in neural information processing systems*, 2000, pp. 1008–1014.
- [46] R. S. Sutton, D. A. McAllester, S. P. Singh, and Y. Mansour, "Policy gradient methods for reinforcement learning with function approximation," in *Advances in neural information processing systems*, 2000, pp. 1057–1063.
- [47] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, and S. Dieleman, "Mastering the game of go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [48] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski et al.,

- “Human-level control through deep reinforcement learning,” *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [49] Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi, “Target-driven visual navigation in indoor scenes using deep reinforcement learning,” in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 3357–3364.
- [50] Z. Ren, X. Wang, N. Zhang, X. Lv, and L. J. Li, “Deep reinforcement learning-based image captioning with embedding reward,” in *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*. IEEE, 2017, pp. 1151–1159.
- [51] J. Li, W. Monroe, A. Ritter, M. Galley, J. Gao, and D. Jurafsky, “Deep reinforcement learning for dialogue generation,” *arXiv preprint arXiv:1606.01541*, 2016.
- [52] X. Zhang and M. Lapata, “Sentence simplification with deep reinforcement learning,” in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2017, pp. 584–594.

APPENDIX

A. Evaluation Metrics

a : candidate sentence, b : set of reference sentences, w_n : n-gram $c_x(y_n)$: count of n-gram y_n in sentence x .

Table IV: Formulation of different metrics.

Metrics	Formulation	Remarks
BLEU	$p_n = \frac{\sum_{w_n \in a} \min \left(c_a(w_n), \max_{j=1, \dots, n } c_{b_j}(w_n) \right)}{\sum_{w_n \in a} c_a(w_n)},$ $BP = \begin{cases} 1 & \text{if } c > r \\ e^{(1-\frac{r}{c})} & \text{if } c \leq r \end{cases},$ $BLEU = BP * \exp \left(\sum_{n=1}^N \alpha_n \log p_n \right).$	r is the reference sentence length, c the length of the candidate sentence, p_n is the n-gram precision of n-grams up to N , α_n is positive weight for each gram.
METEOR	$METEOR = \max_{j=1, \dots, b } \left(\frac{10PR}{R+9P} \right) \left(1 - \frac{1}{2} \left(\frac{\#chunks}{\#mached \text{ unigram}} \right) \right)$	P = unigram precision R = unigram recall chunks is set of unigrams adjacent in a and b_j
ROUGE	$ROUGE = \frac{\sum_{j=1}^{ b } \sum_{w_n \in b_j} \min(c_a(w_n), c_{b_j}(w_n))}{\sum_{j=1}^{ b } \sum_{w_n \in b_j} c_{b_j}(w_n)}$	
CIDER	$CIDER_n = \frac{1}{ b } \sum_{j=1}^{ b } \frac{\mathbf{g}^n(a) \cdot \mathbf{g}^n(b_j)}{\ \mathbf{g}^n(a)\ \cdot \ \mathbf{g}^n(b_j)\ },$ $CIDER = \sum_{n=1}^N \alpha_n CIDER_n$	$\mathbf{g}_n(x)$: vector formed by TF-IDF scores of all n-grams in x , α_n is positive weight for each gram.

Table V: More training details for our model.

Parameter	Value	Description
layers	1	Number of layers in the LSTM encoder/decoder
rnn_size	512	Size of LSTM hidden states
input_feed	1	If value is 1, feed the context vector at each time step as additional input (via concatenation with the word embeddings) to the decoder.
batch_size	64	Batch size
optim	AdaGrad	Optimization method. [sgd—adagrad—adadelta—adam]
lr	1e-3	Initial learning rate
learning_rate_decay	0.5	If update_learning_rate, decay learning rate by this much if (i) perplexity does not decrease on the validation set or (ii) epoch has gone past start_decay_at
start_decay_at	5	Start decaying every epoch after and including this epoch
dropout	0.3	Dropout probability; applied between LSTM stacks.