



Incorporating Domain Knowledge through Task Augmentation for Front-End JavaScript Code Generation

Sijie Shen
Key Lab of High Confidence Software
Technology, MoE (Peking University)
Beijing, China
sjshen@pku.edu.cn

Xiang Zhu
Alibaba Group
Hangzhou, China
hanling.zx@alibaba-inc.com

Yihong Dong
Key Lab of High Confidence Software
Technology, MoE (Peking University)
Beijing, China
dongyh@stu.pku.edu.cn

Qizhi Guo
Alibaba Group
Hangzhou, China
qizhi.gqz@alibaba-inc.com

Yankun Zhen
Alibaba Group
Hangzhou, China
zhenyankun.zyk@alibaba-inc.com

Ge Li*
Key Lab of High Confidence Software
Technology, MoE (Peking University)
Beijing, China
lige@pku.edu.cn

ABSTRACT

Code generation aims to generate a code snippet automatically from natural language descriptions. Generally, the mainstream code generation methods rely on a large amount of paired training data, including both the natural language description and the code. However, in some domain-specific scenarios, building such a large paired corpus for code generation is difficult because there is no directly available pairing data, and a lot of effort is required to manually write the code descriptions to construct a high-quality training dataset. Due to the limited training data, the generation model cannot be well trained and is likely to be overfitting, making the model's performance unsatisfactory for real-world use. To this end, in this paper, we propose a task augmentation method that incorporates domain knowledge into code generation models through auxiliary tasks and a Subtoken-TranX model by extending the original TranX model to support subtoken-level code generation. To verify our proposed approach, we collect a real-world code generation dataset and conduct experiments on it. Our experimental results demonstrate that the subtoken-level TranX model outperforms the original TranX model and the Transformer model on our dataset, and the exact match accuracy of Subtoken-TranX improves significantly by 12.75% with the help of our task augmentation method. The model performance on several code categories has satisfied the requirements for application in industrial systems. Our proposed approach has been adopted by Alibaba's *BizCook* platform. To the best of our knowledge, this is the first domain code generation system adopted in industrial development environments.

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ESEC/FSE '22, November 14–18, 2022, Singapore, Singapore

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9413-0/22/11...\$15.00

<https://doi.org/10.1145/3540250.3558965>

CCS CONCEPTS

• **Software and its engineering** → **Software creation and management**; • **Computing methodologies** → **Artificial intelligence**.

KEYWORDS

Code Generation, Domain Knowledge, Task Augmentation.

ACM Reference Format:

Sijie Shen, Xiang Zhu, Yihong Dong, Qizhi Guo, Yankun Zhen, and Ge Li. 2022. Incorporating Domain Knowledge through Task Augmentation for Front-End JavaScript Code Generation. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '22)*, November 14–18, 2022, Singapore, Singapore. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3540250.3558965>

1 INTRODUCTION

Automatic code generation refers to generating code according to a function description in natural language. The code generation technology can improve the automation level of software development and reduce the workload of software developers, thereby effectively improving the efficiency and quality of software development and maintenance. In recent years, researchers have proposed a series of deep-learning code generation methods with remarkable achievements, such as [5, 18, 19, 25].

Despite the success of the above code generation methods, for all we know, we rarely see the application of code generation in existing industrial systems. Generally, the mainstream code generation methods rely on the accessibility of abundant training data for their success. However, the code in the industrial development environment is usually under a specific domain scenario, where directly paired training data is not always available. In our scenario, we want to generate JavaScript expressions from natural language descriptions, and our major limitation is the insufficient training data problem. In order to perform code generation, we have to build the training dataset from scratch, i.e., mine a group of code and manually write the corresponding semantic descriptions in natural language for each of them. The cost of this process is quite expensive for the construction of a high-quality dataset, so we can only obtain relatively small-scale paired training data. The lack

of training data makes it hard for the neural networks to learn the effective representation to generate code, which dramatically decreases the model performance.

In this paper, we propose a task augmentation method to alleviate the negative effect of insufficient paired training data by incorporating domain knowledge into the code generation model. We design some auxiliary tasks and utilize the domain knowledge that is easily acquired from the codebase and requirement documents as paired training data for the tasks. By training the model together with the code generation main task and auxiliary tasks, the model is able to learn the domain knowledge from the extra tasks and improve its performance on code generation. Furthermore, we present the Subtoken-TranX model by extending the original TranX [25] model to support subtoken-level code generation. The Subtoken-TranX model helps learn the token embedding well and captures the relationship of tokens containing the same subtoken from limited training data. We carry out a series of experiments to demonstrate the effectiveness of task augmentation with domain knowledge and the Subtoken-TranX model. After applying the task augmentation, the Subtoken-TranX achieves a top-1 exact match accuracy of 33.16% and top-5 accuracy of 40.31%, compared to 20.41% and 29.08%, respectively, without task augmentation. The Subtoken-TranX model also outperforms the original TranX model and Transformer model. The results on two categories of data have satisfied the requirements for application in actual industrial systems, and the *BizCook* platform has adopted our approach.

Our contributions can be summarized below:

- We propose a real-world code generation task and the corresponding dataset¹. The dataset consists of 2,489 paired data of JavaScript expressions extracted from Alibaba's codebase and corresponding descriptions in the natural language of Chinese.
- We propose a task augmentation method to incorporate domain knowledge to code generation models through auxiliary tasks.
- We present the Subtoken-TranX model by extending the original TranX model to support subtoken-level code generation.
- We carry out experiments and demonstrate the effectiveness of the task augmentation method and Subtoken-TranX model. Our approach has been adopted by Alibaba's *BizCook* platform and is continuously supporting front-end development.

2 BACKGROUND

The application scenario of our work is the *BizCook* platform in Alibaba. This platform is a front-end development platform tailored based on the business characteristics of *Taobao*. This system covers the main stages of the entire process of front-end development, including requirements, design, coding, and testing. It aims to introduce intelligent approaches to the development process and improve development efficiency. Among them, an important direction of exploration is to generate code from requirement documents and design drafts.

In front-end development, most of the front-end code falls into these three categories:

- User interface code. They are used to control the layout, style, etc., of UI elements, such as the position of a text area and the color of a button.
- Business logic code. They are used to control displayed content of UI elements, such as the text content in a text area or on a button.
- Control flow code. They are used to control the click or other behaviors of elements, such as the on-click event of a button.

Among the three categories, the user interface code is closely related to the design drafts. In our previous works, we have already developed a tool called *imgcook*² to directly generate user interface code from Sketch, Photoshop, and Figma design drafts. On the basis of *imgcook*, we want to extend the code generation ability of the *BizCook* system to other categories. Since the descriptions of business logic code and control flow code are written in the requirement documents, this feature requires the ability to read the natural language descriptions in the requirement documents, convert them into JavaScript expressions, and bind the expressions to UI elements. It is a challenge to generate all types of front-end code simultaneously. Considering that the business logic code accounts for a large proportion of front-end code and is less complicated than the control flow code, we decide first to generate the business logic code at the present stage. This paper presents our practice in dealing with the task of generating the business logic JavaScript code from natural language descriptions.

3 DATASET AND PREPROCESSING

3.1 Code Generation Dataset

In the *BizCook* platform, all the skeletons of code are predefined, and the platform provides ways to set the style, behavior, and display content. During coding, programmers only need to select a UI element and write the target style or string value expressions in JavaScript. The platform will insert the expressions into JavaScript code and bind them to the specific UI element. So all we need to generate is JavaScript expressions. As we stated in the third paragraph of section 2, in the current stage, we only focus on the business logic code expressions that control the display text in the UI elements. Therefore, we randomly collect some expressions from Alibaba's codebase, filter and de-duplicate them, and get a collection of business logic code expressions. In order to get the complete paired data, we still have to manually produce the natural language description under specific rules for each code. Finally, we obtain a paired dataset containing 2,489 examples. We randomly split the dataset into a training set and a test set containing 2293 and 196 examples. When conducting experiments, we randomly sample a subset from the training set as the validation set. Table 1 shows the details of the dataset.

In our dataset, we can further divide the JavaScript logic expressions into four categories:

- string template expression (STE). Code of this category is string expressions produced by filling a string template with variables.

¹Available online at <https://tianchi.aliyun.com/dataset/dataDetail?dataId=107819>.

²<https://www.imgcook.com>.

Table 1: Details of code generation dataset. Since the description is written in Chinese, the token numbers are counted as the number of Chinese characters. All details are counted after the preprocessing described in section 3.2.

	Train Set	Test Set
# Examples	2,293	196
# Desc. Tokens	43,719	3,874
# Code Tokens	25,806	2,256
# Avg. Desc. Tokens	19.07	19.77
# Avg. Code Tokens	11.25	11.51

- OR logic expression (OLE). Code of this category are several values joined by short circuit OR operator.
- condition expression (CE). Code of this category is often a ternary conditional expression.
- data processing expression (DPE). Code of this category often contains processing to a data, such as taking a substring of a string.

Table 2 shows examples of each category. In actual use, the input descriptions are in Chinese, and we provide the corresponding English translation here. We tag every piece of data in the test set with category labels and count the test set details by category. Table 3 shows the details of test set by category.

3.2 Preprocessing

Considering that our dataset scale is small and the variability and noise in the dataset can affect code generation performance, we perform a series of preprocessing on the original data. In this section, we introduce three main preprocessing methods we use.

3.2.1 Code Canonicalization. In our dataset, the code have some stylistic differences since the codes are not written by the same person. For example, developers have different preferences of using single and double quotes in string literals and whether to add semi-colons at the end of statements, etc. We use the Esprima³ parser to parse the JavaScript code into an abstract syntax tree (AST) and then use the Escodegen⁴ code generator to convert the AST back to a canonicalized code. In this way, we can eliminate the stylistic differences between the codes to a certain extent.

3.2.2 String Literal Replacement. In our dataset, a proportion of code contains string literals. According to our regulations, all the string literals in the code must appear in the natural language description. Therefore, we can use placeholders to replace these string literals to simplify the code generation. As shown in the table 4, for string literals that do not contain variables, we use placeholders such as <STR1>, <STR2> to replace string literals. For string literals containing variables such as “xx”, “yy”, etc., we replace these string literals segment-wise with placeholders. These placeholders will be considered a single token by the model.

3.2.3 Member Access Simplification. A lot of code includes access to member variables. We simplified the member variable access in the code by removing the accessed object and only keeping the

accessed fields. For examples, `task.status` is simplified to `status` and `task.assets.completeBtn` is simplified to `completeBtn`. We simplify the member access because we cannot predict the accessed object correctly from the description without programming context. Other modules in the development platform will search the programming context for the proper accessed object using techniques such as code static analysis and will complete the member access expression.

4 METHODOLOGY

We explain the details of our code generation method in this section. In section 4.1, we present our method of applying a variable semantic table for task augmentation. In section 4.2, we describe the architecture of our Subtoken-TranX model. In section 4.3, we describe the details of how we make Subtoken-TranX support JavaScript language, which also applies in the original TranX model.

4.1 Task Augmentation

As we face insufficient training data, we want to leverage external domain knowledge to assist with code generation. We found that variable names occupy a large part of the front-end JavaScript code we study. Thus, it benefits a lot if we make the model learn more about the correct use of variable names from the domain knowledge. Following this idea, we think it is helpful if we can obtain a paired dataset of variable names and their semantic meaning and incorporate this data into the code generation model. We extract a variable semantic table containing the name and semantic description of variables commonly used in front-end development. We collect these variable–description paired data from various sources, including the following:

- Requirement documents. The requirement documents may contain some variable name conventions for variables used across modules in the project. We can gather these variable name conventions and add them to the variable semantic table.
- Codebase. In the development of large-scale projects, developers usually use some protocols for serializing data structures (Protocol Buffer, Thrift, etc.) to describe different data structures. For the fields in the data structures, there are usually comments that describe the meanings of the fields. We can collect these fields and semantic meanings and add them to the variable semantic table.
- Database field definitions. Tables in the database may have corresponding documents to describe the semantics of each field in the table. We can also use these fields and semantic meanings to build the variable semantic table.

Eventually, we collect a variable semantic table containing 15,525 variables that are commonly used in front-end developments. Some of the examples are shown in table 5.

There are many ways to incorporate the variable semantic table into the code generation model. For example, we can use the variable semantic table to pre-train the model and then use the code generation data to fine-tune the model. In our attempts, we use a task augmentation method to leverage the variable semantic table for code generation.

³<https://esprima.org>

⁴<https://github.com/esutils/escodegen>

Table 2: Examples of data in different categories.

Category	String template expression
Description	显示“优惠券已抵扣xx元”，xx为折扣价
Desc. Trans.	Displays “The coupon has been deducted xx yuan”, xx is the discounted price
Code	{`优惠券已抵扣\${discountPrice}元`}
Category	OR logic expression
Description	动态展示用户昵称，兜底为空
Desc. Trans.	Dynamically display the user’s nickname, the default value is empty
Code	{ user && user.nick " " }
Category	Condition expression
Description	如果内容类型为直播，则展示直播时间描述，否则展示营销时间描述
Desc. Trans.	If the content type is live, show the description of the live time, otherwise show the description of the marketing time
Code	{contentType === 'live' ? liveTimeDesc : marketingTimeDesc}
Category	Data processing expression
Description	动态展示金币展示价格小数部分
Desc. Trans.	Dynamically display the fractional part of coin show price
Code	{`\${data.coinShowPrice.split(".")[1]}`}

Table 3: Details of test data by category.

	ST	OLE	CE	DPE
# Examples	49	71	71	5
# Desc. Tokens	662	1,242	1,856	114
# Code Tokens	503	609	1,053	91
# Avg. Desc. Tokens	13.51	17.49	26.14	22.80
# Avg. Code Tokens	10.27	8.58	14.83	18.20

Task augmentation is the way of training a neural network with multiple tasks. In that way, the model can learn more knowledge from the auxiliary tasks and improve the main task’s performance. In our practice, we take the variable semantic table as domain knowledge and design auxiliary tasks to train the model with the main code generation task. This task augmentation method has two main benefits. First, the model can learn more about the relations between variable semantics and variable names from the auxiliary task and use it in code generation. Thus the model may make better predictions of variable names during code generation and improve the overall performance. Second, the addition of auxiliary tasks enlarges the training data size and reduces the overfitting of the model.

Many different auxiliary tasks can leverage the variable semantic table as training data. The most straightforward task is to predict variable names from variable semantic meanings. This task requires the model to understand the variable semantic meanings and predict the correct variable name. However, the output of this auxiliary task is not a complete code, and naturally, we can not convert it into a legal abstract syntax tree. So we cannot apply this auxiliary task to tree-based code generation methods such as TranX. We can only use this auxiliary task on token sequence-based methods.

In order to apply the task augmentation in AST-based code generation methods, we propose another auxiliary task. We found that for a string template expression that has no other string literals to join with, the code is just a variable name plus syntax symbols such as braces and semicolons. We can transform the variable semantic

table into these string template expressions. Taking the “picUrl-图片链接(link of the picture)” as an example, we can rewrite the input as “展示图片链接(show link of the picture)” and the output as “{ picUrl; }”. In that way, we write the output to a legal code and can be parsed to an abstract syntax tree so that we can apply the auxiliary task to AST-based code generation methods. Through that rewrite, we have also largely narrowed the gap between the main task and auxiliary task, enabling task augmentation methods to be more effective. We adopt this auxiliary task in our code generation model. We will compare this auxiliary task with the variable names predicting task for token sequence-based code generation on the Transformer model. We’ll also compare the task augmentation method with the pre-train method.

4.2 Subtoken-TranX Model

The original TranX model generates code at the token level. In that way, an identifier corresponds to a token in the vocabulary. This setting is not friendly to code generation under small-scale training data. The token-level vocabulary is usually large, so the model needs to maintain a large embedding matrix, which makes the model easy to overfit. And the identifiers in code may appear only a few times in the training set, so it is hard for the model to learn a good word embedding representation. Besides, the model may fail to capture the relationship of tokens containing the same subtoken from limited training data. A better way is to generate code at the subtoken level.

To generate code at the subtoken level, we subtokenize all the identifiers according to case boundaries for identifiers in the camel case and underscores for identifiers in the snake case [2] for all the code. For example, the identifier liveTimeDesc will be split to live, ##Time, and ##Desc. The prefix ## indicates that this subtoken and the previous one are split from the same token, and we need to join them together when converting subtokens back to tokens. We choose not to use statistical-based subtokenization methods such as BPE [17] because we think the distribution of small-scale corpus is not statistically significant. Thus, the subtokens learned with

Table 4: Examples of string literal replacement

Original Description	判断是否幸运, 条件成立则显示‘恭喜你押中啦’, 否则显示‘很遗憾未押中’
Ori. Desc. Translation	Test if lucky or not, display ‘Congratulations on your bet’ if the conditions are met, otherwise display ‘Unfortunately not betting’
Original Code	{isLucky ? ‘恭喜你押中啦’ : ‘很遗憾未押中’}
Simplified Description	判断是否幸运, 条件成立则显示‘<STR1>’, 否则显示‘<STR2>’
Simp. Desc. Translation	Test if lucky or not, display ‘<STR1>’ if the conditions are met, otherwise display ‘<STR2>’
Simplified Code	{isLucky ? ‘<STR1>’ : ‘<STR2>’}
Original Description	显示‘满xx使用’, xx为起步费
Ori. Desc. Translation	Displays ‘For orders over xx, use the coupon’, xx is the starting fee
Original Code	{‘满’ + startFee + ‘使用’}
Simplified Description	显示‘<STR1> xx <STR2>’, xx为起步费
Simp. Desc. Translation	Displays ‘<STR1> xx <STR2>’, xx is the starting fee
Simplified Code	{‘<STR1>’ + startFee + ‘<STR2>’}

Table 5: Examples of the variable semantic table

Variable Name	Variable Semantic
shopLogo	店铺标志(shop logo)
beforePromotionPrice	促销前价格(price before promotion)
storeCityName	门店所属城市(city the store located)
roomStatus	直播状态(live room status)
picUrl	图片链接(link of the picture)
...	...

BEP from the small-scale corpus may not truly reflect the actual distribution of subtokens and may not achieve the best performance.

To perform the subtoken-level code generation, we build our Subtoken-TranX model based on the original TranX model. There are three stages to transition from natural language description to code in Subtoken-TranX. First, the encoder-decoder neural network in TranX takes the natural language description as input and then generates a sequence of actions to construct the AST. In this stage, the model will check the syntax of the target programming language and make sure the generated actions can be used to build a valid AST. Second, we construct the AST with the generated action sequence in the first stage. Third, we convert the AST to target code. The AST-constructing actions defined in the original TranX model have three types:

- **APPLYCONSTR**[c]. This action applies a construction rule c on current field of the AST under construction.
- **REDUCE**. This action marks the end of generation of the current field with optional or multiple cardinalities.
- **GENTOKEN**[v]. This action generates a terminal token as a leaf node to the current field.

In the original TranX model, the generation of a terminal token corresponds to one and only one **GENTOKEN** action, so the original TranX model only supports token-level code generation but not subtoken-level code generation. To enable the model to support subtoken-level generation, we need to do some modifications to the **GENTOKEN** action. We replace the **GENTOKEN** action with **GENSUBTOKEN** in the Subtoken-TranX model. We stipulate that at each position of AST that needs to generate a terminal symbol,

the model can generate multiple **GENSUBTOKEN** actions, of which each **GENSUBTOKEN** action generates a subtoken. We use a special token $\langle \text{EOT} \rangle$ to indicate that all the subtokens of the token in the current field have been completely generated. The subtokens generated by these consecutive **GENSUBTOKEN** actions will be joined together to form the token generated at the current position and inserted into the AST under construction. The model will then proceed to the generation of the next field. With this modification, the Subtoken-TranX model can support subtoken-level code generation. Figure 1 show an example of JavaScript AST and the sequence of actions used to construct the AST. In this example, the identifier `contentType` contains two subtokens which correspond to the **GENSUBTOKEN** actions in time step $t_{7,1}$ and $t_{7,2}$. In time step $t_{7,3}$, the model generates the $\langle \text{EOT} \rangle$, then it merges all the consecutive subtokens generated in this field to the token `contentType` and inserts it to the AST under construction. The generation of `liveTimeDesc` and `marketingTimeDesc` is also similar.

Noticing that this modification is target programming language agnostic, the Subtoken-TranX model can be used to generate code in other programming languages besides JavaScript. We can replace the original TranX model with the Subtoken-TranX model whenever subtoken-level generation is preferred.

The architecture of neural networks in Subtoken-TranX follows the original TranX in [25]. We compute the probabilities of generating an action sequence \mathbf{z} as

$$p(\mathbf{z}|\mathbf{x}) = \prod_t p(a_t | a_{<t}, \mathbf{x}). \quad (1)$$

The encoder is a bidirectional LSTM [8] which encodes the input utterance $\{x_i\}_{i=1}^n$ into representations $\{\mathbf{h}_i\}_{i=1}^n$. The decoder is also an LSTM network which computes the hidden state \mathbf{s}_t at each time step as

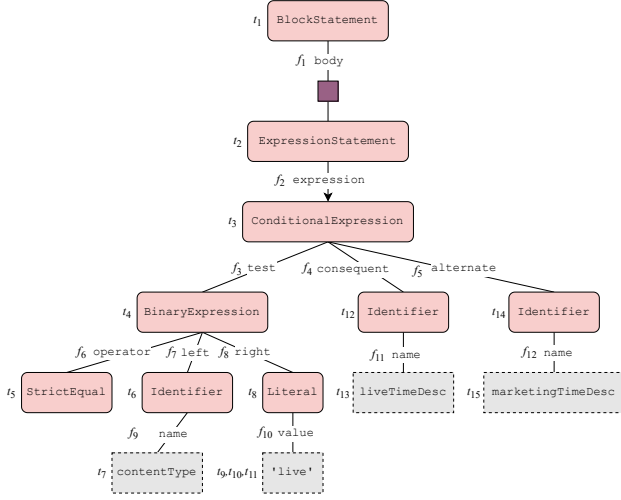
$$\mathbf{s}_t = f_{\text{LSTM}}([\mathbf{a}_{t-1}; \tilde{\mathbf{s}}_{t-1}; \mathbf{p}_t], \mathbf{s}_{t-1}). \quad (2)$$

The inputs of decoder LSTM are three vectors. \mathbf{a}_{t-1} is the embedding of last action. $\tilde{\mathbf{s}}_{t-1}$ is the attention result computed as

$$\tilde{\mathbf{s}}_t = \tanh(W_c [c_t; \mathbf{s}_t]), \quad (3)$$

where c_t is the context vector computed in attention [14] with encoded results $\mathbf{h}_{i=1}^n$. \mathbf{p}_t is parent feeding information which is the concatenation of the embedding of the frontier field \mathbf{n}_{f_t} and s_{p_t} ,

```
Code: {contentType === 'live' ? liveTimeDesc : marketingTimeDesc}
```



t	n_{f_t}	Action
t_1	root	BlockStatement(stmt* body)
t_2	f_1	ExpressionStatement(expr expression)
t_3	f_2	ConditionalExpression(expr test, expr alternate, expr consequent)
t_4	f_3	BinaryExpression(binary_operator operator, expr left, expr right)
t_5	f_6	StrictEqual()
t_6	f_7	Identifier(name)
$t_{7,1}$	f_9	GENSUBTOKEN[contentType]
$t_{7,2}$	f_9	GENSUBTOKEN[##Type]
$t_{7,3}$	f_9	GENSUBTOKEN[<EOT>]
t_8	f_8	Literal(literal? value)
t_9	f_{10}	GENSUBTOKEN[<SOS>]
t_{10}	f_{10}	GENSUBTOKEN[live]
t_{11}	f_{10}	GENSUBTOKEN[<EOS>]
t_{12}	f_4	Identifier(identifier name)
$t_{13,1}$	f_{11}	GENSUBTOKEN[live]
$t_{13,2}$	f_{11}	GENSUBTOKEN[##Time]
$t_{13,3}$	f_{11}	GENSUBTOKEN[##Desc]
$t_{13,4}$	f_{11}	GENSUBTOKEN[<EOT>]
t_{14}	f_5	Identifier(identifier name)
$t_{15,1}$	f_{12}	GENSUBTOKEN[marketing]
$t_{15,2}$	f_{12}	GENSUBTOKEN[##Time]
$t_{15,3}$	f_{12}	GENSUBTOKEN[##Desc]
$t_{15,4}$	f_{12}	GENSUBTOKEN[<EOT>]
t_{16}	f_1	REDUCE (close the frontier field f_1)

Figure 1: An example of JavaScript AST and the action sequence to generate the AST. In this example, the generation of token `liveTimeDesc` of field f_{11} and `marketingTimeDesc` of field f_{12} is split into several actions with each action generates a subtoken.

the decoder’s state at which the constructor of n_{f_t} is generated by the `APPLYCONSTR` action.

We compute the probability of action `APPLYCONSTR`[c] as

$$p(a_t = \text{APPLYCONSTR}[c] \mid a_{<t}, \mathbf{x}) = \text{softmax}(\mathbf{a}_c^\top \mathbf{W}\tilde{\mathbf{s}}_t). \quad (4)$$

The probability of `GENSUBTOKEN` is a hybrid probability of generation and copy from input, which is formulated as

$$p(a_t = \text{GENSUBTOKEN}[v] \mid a_{<t}, \mathbf{x}) = p(\text{gen} \mid a_t, \mathbf{x})p(v \mid \text{gen}, a_t, \mathbf{x}) + p(\text{copy} \mid a_t, \mathbf{x})p(v \mid \text{copy}, a_t, \mathbf{x}), \quad (5)$$

where

$$p(v \mid \text{gen}, a_t, \mathbf{x}) = \text{softmax}(\mathbf{a}_c^\top \mathbf{W}\tilde{\mathbf{s}}_t), \quad (6)$$

$$p(v \mid \text{copy}, a_t, \mathbf{x}) = \text{softmax}(\mathbf{h}_t^\top \mathbf{W}\tilde{\mathbf{s}}_t), \quad (7)$$

and

$$p(\text{gen} \mid \cdot), p(\text{copy} \mid \cdot) = \text{softmax}(\mathbf{W}\tilde{\mathbf{s}}_t). \quad (8)$$

We add beam search on the model so that the model will predict top- k results for every input utterance.

4.3 JavaScript Language Support

The TranX model (and our Subtoken-TranX model) uses ASDL [21] to describe the syntax of programming languages. Most of the previous works that use the TranX model for general-purpose programming language code generation [9, 23, 25] use Python as the target language. The Python language officially uses ASDL to describe the syntax rules. So it is convenient to utilize the TranX model for Python code generation. For the Subtoken-TranX model to support the generation of the JavaScript code, we need to construct the ASDL for the syntax rules of the JavaScript language.

ASDL is a relatively simple abstract syntax description language. It only supports using exactly one type to constrain a field and does not support using multiple types to constrain a field. It also does not support using subtype to constrain a field. For programming languages with complex syntax rules such as JavaScript, we can only

support parts of the syntax rules of JavaScript and make particular adaptations for some of the syntaxes.

Table 6 show 2 examples of JavaScript’s abstract syntax that are incompatible with ASDL. In the `CallExpression` of JavaScript, the type of field `callee` can be either `Expression` or `Import`. Given that the ASDL only supports using exactly one type to constrain a field, we just keep the `Expression` type for the field and remove support for `Import` type for mostly the callee fields are of type `Expression`. We do this kind of choice for all the syntaxes with multiple types for a field and keep only the most frequently used type. In `BreakStatement`, the `label` field is specified as `Identifier`, which is not a type but a subtype of `Statement` (or `stmt` in our ASDL since we use lower case abbreviations to represent types). Now that ASDL does not support using subtype to constrain a field, we directly use its parent type `stmt` as the constrain instead of `Identifier` to comply with the rules of ASDL. We relax the type constraints in this way for all syntaxes that has the subtype issue. This treatment brings up another problem that the generated AST containing these syntaxes may comply with ASDL syntax constraints but is not a legal JavaScript AST. If we encounter this situation during code generation, we simply discard this result and adopt other results predicted by the model.

We write the ASDL of JavaScript regarding the JavaScript abstract syntax in the form of Mozilla Parser API provided in the documentation of Esprima parser. The ASDL of JavaScript we write can cover the vast majority of the front-end JavaScript code in our study scenario.

In addition to the ASDL of JavaScript, we write two functions to do the conversion between ASDL AST and JavaScript AST. The functions traverse the input AST and build the output AST according to the information of AST nodes. We use Esprima and Escodegen to do the conversion between JavaScript AST and JavaScript code.

Table 6: Examples of ASDL-incompatible JavaScript abstract syntax

```
Syntax of CallExpression:
interface CallExpression {
  type: 'CallExpression';
  callee: Expression | Import;
  arguments: ArgumentListElement[];
}

ASDL of CallExpression:
expr = CallExpression(expr callee, expr* arguments)
```

```
Syntax of BreakStatement:

interface BreakStatement {
  type: 'BreakStatement';
  label: Identifier | null;
}

ASDL of BreakStatement:
stmt = BreakStatement(expr? label)
```

So far, we have added JavaScript language support to the Subtoken-TranX model and can use it to generate JavaScript code from natural language descriptions.

5 EXPERIMENTAL SETTINGS

5.1 Setup

5.1.1 Dataset. We use the dataset presented above in section 3 for our experiments. We randomly sample a subset from the training set as our validation set. In the task augmentation, we use the variable semantic table of 15,525 variables that we collected from the codebase, database field definitions, and requirement documents.

5.1.2 Baselines. Considering that there are two main categories of code generation models, tree-based and sequence-based, we selected a typical model from each category as our baselines. They are:

- **TranX.** The TranX model is an AST-based token-level code generation model and performed well in previous work.
- **Transformer.** The Transformer model is widely-used and achieves good performance in many sequence-to-sequence tasks. We use it for subtoken-level code generation.

We compare our proposed Subtoken-TranX model with the two baselines, and we also evaluate our task augmentation method on the Subtoken-TranX and the baseline models.

5.1.3 Implementation Details. For the Subtoken-TranX and baseline TranX model, we use the tools Esprima and Esgen for the conversion between JavaScript code and AST. We also use Esprima to tokenize the JavaScript code for the baseline Transformer. We tokenize the natural language description with NLTK and Jieba. We subtokenize the identifier according to the case boundaries for identifiers in the camel case and the underscores for those in the

snake case. The Subtoken-TranX model is built based on the open-source code of TranX⁵, and the Transformer model is built with the official implementation of the Transformer model in PyTorch.

For the Subtoken-TranX and TranX models, we set the hidden size of the encoder and decoder LSTM to 256. The embedding size of subtokens and actions is 128. During the training process, we use a batch size of 32 and train the model for 300 epochs. We use Adam [10] optimizer and set the initial learning rate to 1×10^{-3} .

For the Transformer model, we use the encoder and decoder with 4 layers. We set the model size to 128, and the size of the feed-forward network to 512. Each multi-head attention in the Transformer model has 4 heads. During training, we use a batch size of 32 and train the model for 300 epochs. We use the AdamW [12] optimizer to optimize the model parameters. We use the learning rate warming up and decay during the training process and set the max learning rate to 1×10^{-4} .

We add beam search on both Subtoken-TranX and Transformer and set the beam width to 5 for both models.

5.2 Evaluation Metrics

We use the following evaluation metrics for the code generation:

- **Exact Match Accuracy.** The exact match accuracy regards a predicted result as true when the result is exactly the same as the reference code. This is a very strict metric. For codes that are functionally identical but literally different, the exact match accuracy regards the prediction as false.
- **BLEU [15].** The BLEU score checks the precision of n-grams in the predicted result and computes a composite score. We use the `corpus_bleu` in NLTK to compute the BLEU score of our result.
- **Edit Similarity.** The edit similarity computes the similarity of two strings with Levenshtein edit distance. When used in code generation, this metric indicates how much modification the user needs to change the generated code to the correct reference code.

6 EXPERIMENTAL RESULTS

6.1 Main Results

To verify the effectiveness of our task augmentation method that leverages external domain knowledge for code generation and our Subtoken-TranX model, we train the models with and without task augmentation and compare the results of different models on code generation. Table 7 show the main results of code generation. The results with “+TA” are the results when training the models with task augmentation. The results show that the Subtoken-TranX outperforms the original TranX and Transformer model with or without the task augmentation applied. When applied with task augmentation, both the Transformer and the Subtoken-TranX model that generate code at the subtoken level have significantly improved their performance. The top-1 accuracy and top-5 accuracy have improved by more than 10%. The results demonstrate the effectiveness of our task augmentation method that incorporates external domain knowledge for code generation on subtoken-level generation models. We also notice that the results of TranX decrease after

⁵<https://github.com/pcyin/tranX>

Table 7: Main results of code generation

Method	Acc-1	Acc-5	BLEU	EditSim
Transformer	16.84	26.02	67.98	77.31
TranX	17.35	22.45	67.56	76.20
Subtoken-TranX	20.41	29.08	66.77	77.63
Transformer+TA	29.08	39.29	69.72	82.92
TranX+TA	16.33	27.55	61.66	74.77
Subtoken-TranX+TA	33.16	40.31	71.94	85.27

Table 8: Results of variable usage in code generation.

Method	Precision	Recall	F1
Transformer	27.66	27.27	27.46
TranX	26.67	27.97	27.30
Subtoken-TranX	30.57	33.57	32.00
Transformer+TA	41.99	41.26	41.62
TranX+TA	36.12	37.76	36.92
Subtoken-TranX+TA	50.00	48.25	49.11

applying the task augmentation. We think this may be because the addition of the semantic mapping table only has a limited effect on the token-level generation model. On the contrary, the increase in the vocabulary size will increase the number of models and make it easier to overfit so that the model performance will decrease.

6.2 Results of Variable Usage

We conjecture that the task augmentation with the variable semantic table enables the model to learn more about the relationship between variable semantics and variable names from auxiliary tasks, thereby improving the accuracy of variable usage during code generation. To verify our conjecture, we verify the results of variable usage in code generation. We extract all variable names in the predicted results and the reference code on the test set and calculate the precision, recall, and F1 of the variable usage. Table 8 show the results of variable usage in code generation. We can see that the metrics of variable usage have great improvement on all the models after applying the task augmentation. The results validate our conjecture that the addition of auxiliary tasks can improve the accuracy of variable usage in code generation. Among all the models, the Subtoken-TranX model achieves the best precision, recall, and F1 for variable usage.

6.3 Results of Different Categories

We calculated the metrics of methods on different categories of test data. As mentioned above, we divide the data into four categories: string template expression, OR logic expression, condition expression, and data processing expression. Figure 2 shows the result of the four categories. We only show the results of Subtoken-TranX and Transformer, as the TranX and Subtoken-TranX models share a similar architecture, and the latter one performs better.

Due to the category of data processing expression only having 5 examples in the test set, results on this category of data are not statistically significant, so we focus more on the results of

the other three categories. Overall, the performance of the string template expressions and OR logic expressions is better than that of the condition expressions and data processing expressions. This is because the difficulty of generating condition expressions and data processing expressions condition expressions is greater than that of the other two categories. The condition expressions and data processing expressions may involve more complex logic, and the average number of tokens is significantly longer. Besides, there are only a few data processing expressions for training, so the generation of this category is even harder.

We can see from the results that after applying the task augmentation method, the performance of models have different magnitudes of improvement on various categories of data. The results also show that the Subtoken-TranX model outperforms the Transformer model on three categories of data except OR logic expressions. This inconsistency may be because the natural language descriptions of condition expressions have more complex semantics, and the stronger representation ability of the Transformer model can better capture the semantics in natural language descriptions, thereby generating better results.

For the string template expressions, the top-5 accuracy of Subtoken+TA is 44.90%, and the editing similarity reaches 86.21%. For the OR logic expressions, the top-5 accuracy of Transformer+TA is 52.11%, and the editing similarity reached 88.40%. The model performance on these two code categories has met the requirements for application on actual industrial systems. Our research achievements have been adopted to Alibaba’s *BizCook* system.

6.4 Results of Different Usages of Variable Semantic Table

In section 4.1, we mentioned that there are many different ways to leverage the variable semantic table for code generation. A straightforward approach in the task augmentation method is to use it for an auxiliary task that predicts variable names from variable semantic meanings. Even though we cannot use this auxiliary task on AST-based code generation methods such as Subtoken-TranX, we can use it on token sequence-based methods such as Transformer. We compare the code generation results using different auxiliary tasks for task augmentation. Besides the task augmentation method, we can also use the variable semantic table for pre-training and fine-tune the model with the code generation dataset.

Table 9 show the code generation results of incorporating the variable semantic table into the model in different ways. The line of Transformer is the result of Transformer model training only with code generation dataset. The “+PT” is the result of pre-training model with variable semantic table and fine-tuning with code generation dataset. The “+VP” is the result of training model with the auxiliary task of variable name prediction from variable semantics. The “+CG” is the result of training model with the auxiliary task after rewriting the variable semantic table into code generation paired data. The results show that all the methods of using the variable semantic table are effective in improving code generation performance. In line with our expectation, the task augmentation method with the code generation auxiliary task (+CG) achieves the best performance. It outperforms the variable prediction task augmentation (+VP) because the auxiliary task and the main task

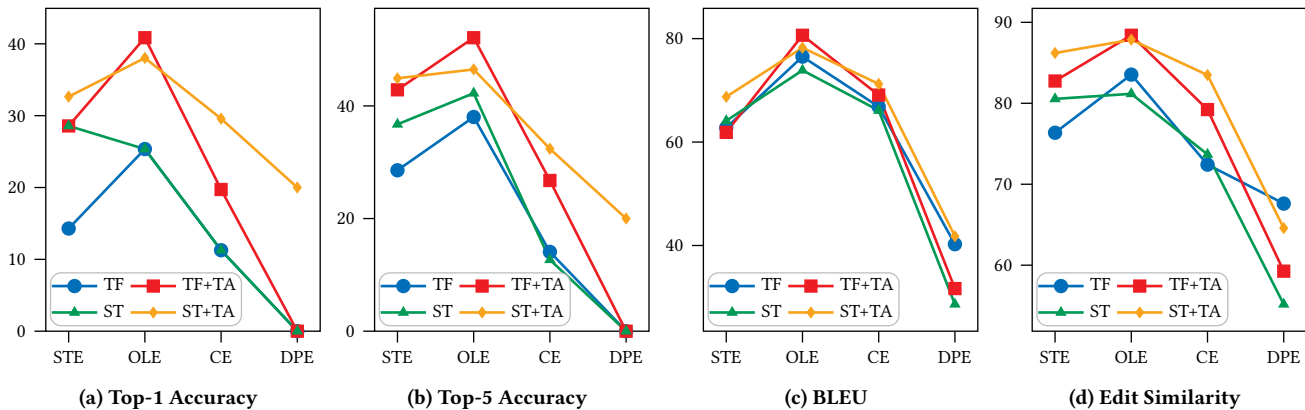


Figure 2: Metrics of different categories on test data

Table 9: Results of different uses of variable semantic table

Method	Acc-1	Acc-5	BLEU	EditSim
Transformer	16.84	26.02	67.98	77.31
Transformer+PT	26.53	38.27	69.91	82.47
Transformer+VP	25.00	37.24	68.82	82.38
Transformer+CG	29.08	39.29	69.72	82.92

have a smaller gap. Similarly, in the pre-train method (+PT), since the model only needs to learn one task in the fine-tuning stage, this method also performs better than variable prediction task augmentation (+VP).

6.5 Case Study

We conducted a sample analysis to compare the generated results of different methods visually. We picked two individual examples from the test dataset to show the generation results of models.

Table 10 shows the first example. We can see from the results that both the Subtoken-TranX and Transformer models failed to predict the variable `trainHeadTitle` without applying the task augmentation method. The models may capture some related information from the natural language and generate corresponding subtokens such as `head` and `title`. But finally, they fail to combine all the subtokens in the correct way to generate the correct variable name. We examine the variable semantic table we extracted and found the variable `trainHeadTitle` exists in the table. After applying the task augmentation, the model can learn the correct variable name from the auxiliary tasks. Thus, both Subtoken-TranX and Transformer generate the correct result in the top-1 prediction.

Table 11 shows the second example. We can find that in the Transformer results without task augmentation, most of the results contain the variable `subTitle` and the number 16. We found an example in training data that is very similar to this example. Just replacing the `title` to `subTitle` and 15 to 16 in this test example, we get the similar code in training data, which is exactly the same with the top-1 prediction of Transformer in this test example. This behavior implies that the Transformer model has large

overfitting. The model "remembers" the content of the training data set and directly outputs the same results when encountering similar samples. After applying task augmentation, the top-2 result of Transformer+TA and the top-1 result of Subtoken-TranX+TA gives the correct predictions. This improvement shows that we have alleviated the over-fitting of the model to a certain extent by expanding the size of training data with task augmentation and improving the model's performance.

7 RELATED WORK

This paper is related to deep learning-based code generation methods and how to incorporate external knowledge into code generation.

Code generation usually refers to generating code according to a natural language description. Some works also refer to statement-level code completion as code generation [20]. In natural language processing, the semantic parsing task that converts natural language to a logical form is similar to code generation. The solutions for these two tasks are basically the same. So we make no distinction between related work on these two tasks and call them code generation collectively. Overall, code generation methods with neural networks can be divided into two categories that are token-based methods and tree-based methods. The token-based methods regard the input utterance and output code as token sequences and use seq2seq network [4, 11] or Transformer to generate the output sequence from the input sequence. The tree-based methods have various attempts. The seq2tree network [4] generates a tree-structured representation in top-down and breadth-first order from natural language input. The Abstract Syntax Networks [16] generate an abstract syntax tree in depth-first order by selecting and applying a constructor to the AST under construction. That makes the output tree conform to the programming language syntax. The TranX [25] generates a sequence of actions that construct an AST. It also utilizes the programming language syntax to filter the actions that do not meet the syntax. The model architecture is more simple than ASN and also achieves good performance. Some other improvements to the TranX model have also resulted in improved performance [9, 23].

Table 10: Case study example 1.

Description	动态展示火车头标题，兜底显示‘春运火车票’
Desc. Translation	Display the title of the train head, the default value is “Spring Festival Train Ticket”.
Reference Code	{trainHeadTitle ‘春运火车票’;}
Subtoken-TranX Result	{downTitle ‘春运火车票’;}
Transformer Result	{headBannerTitle ‘春运火车票’;}
Subtoken-TranX+TA Result	{trainHeadTitle ‘春运火车票’;} ✓
Transformer+TA Result	{trainHeadTitle ‘春运火车票’;} ✓

Table 11: Case study example 2.

Description	判断是否显示中间，成立则展示标题前15个字，否则展示前10个字或者不展示
Desc. Translation	If show the middle , display the first 15 words of the title, otherwise display the first 10 words or not display
Reference Code	{isShowMid ? title.substring(0, 15) : title.substring(0, 10) “”;
Subtoken-TranX Result	R1. {isShowMid ? title.substring(0, null) : title.substring(0, 10) null;} R2. {isShowMid ? title.substring(0, 15) : title.substring(0, 10) null;}
Transformer Result	R1. {isShowMid ? subTitle.substring(0, 16) : subTitle.substring(0, 10) “”}; R2. {isShowSuccess ? subTitle.substring(0, 16) : subTitle.substring(0, 10) “”};
Subtoken-TranX+TA Result	R1. {isShowMid ? title.substring(0, 15) : title.substring(0, 10) “”;} ✓ R2. {isShowMid ? title.substring(0, 15) : title.substring(null, 15) “”};
Transformer+TA Result	R1. {isShowMid ? title.substring(0, 16) : title.substring(0, 10) “”} R2. {isShowMid ? title.substring(0, 15) : title.substring(0, 10) “”} ✓

In addition to innovations in model architectures, researchers also have proposed many ways to incorporate external knowledge into code generation. There are retrieval-based methods that search for a similar code snippet with the input natural language. The code generation model can use the retrieved code to guide the code generation [7, 26] or predict the output by editing the retrieved code [6]. Some works utilize text–code paired corpus mined from Stack Overflow as an external knowledge to pre-train the code generation model [24]. With the widespread application of pre-training techniques, many works leverage models pre-trained with large-scale corpus data for code generation [1, 3, 13, 22] and achieve remarkable performance.

8 CONCLUSION AND FUTURE WORKS

This paper presents our practice of applying the JavaScript code generation method for front-end development on Alibaba’s *BizCook* platform. Facing the problem of insufficient training data, we incorporate external domain knowledge for code generation through task augmentation so that the code generation model can perform better with the help of auxiliary training data. We also extend the TranX model and build a Subtoken-TranX model for subtoken-level code generation, which is also helpful in the case of limited training data. We carry out a series of experiments and demonstrate the effectiveness of our methods compared to baseline methods in the case of limited training data. The results show that our code generation method has met the application requirements on actual industrial systems in several code categories and has been adopted to Alibaba’s *BizCook* system for production.

So far, we can see that the model does not perform as well on condition expressions and data processing expressions as the other two categories. Our current task augmentation method only takes a variable semantic table as domain knowledge. It thus improves the accuracy of variable usage, which is most effective for codes that are sensitive to variable names. For condition expressions, the variable semantic table cannot improve the prediction of conditions, which is an essential part of condition expressions. For data processing expressions, the performance heavily relies on the prediction of data processing API, which is also not presented in the variable semantic table. In future works, we can mine a series of conditions and their natural language descriptions from the codebase and a series of data processing APIs from documents. We can use a similar method as presented in this paper to design auxiliary tasks and leverage these data for task augmentation. In that way, we can improve the code generation performance on condition expressions and data processing expressions. In addition, with the launch of the code generation tool, we can continuously collect the usage history as paired data and use it for model training. The increase in the amount of data will also help improve the performance of the model on various categories of data.

ACKNOWLEDGMENTS

This research is supported by Alibaba Group through Alibaba Innovative Research Program and the National Natural Science Foundation of China under Grant No. 62072007, 62192733, 61832009, 62192731, 62192730.

REFERENCES

- [1] Wasi Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. 2021. Unified Pre-training for Program Understanding and Generation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 2655–2668.
- [2] Dave Binkley, Marcia Davis, Dawn Lawrie, and Christopher Morrell. 2009. To camelcase or under_score. In *2009 IEEE 17th International Conference on Program Comprehension*. IEEE, 158–167.
- [3] Colin Clement, Dawn Drain, Jonathan Timcheck, Alexey Svyatkovskiy, and Neel Sundaresan. 2020. PyMT5: multi-mode translation of natural language and Python code with transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 9052–9065.
- [4] Li Dong and Mirella Lapata. 2016. Language to Logical Form with Neural Attention. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 33–43.
- [5] Li Dong and Mirella Lapata. 2018. Coarse-to-Fine Decoding for Neural Semantic Parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 731–742.
- [6] Tatsunori B Hashimoto, Kelvin Guu, Yonatan Oren, and Percy S Liang. 2018. A retrieve-and-edit framework for predicting structured outputs. *Advances in Neural Information Processing Systems* 31 (2018).
- [7] Shirley Anugrah Hayati, Raphael Olivier, Pravalika Avvaru, Pengcheng Yin, Anthony Tomic, and Graham Neubig. 2018. Retrieval-Based Neural Code Generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 925–930.
- [8] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [9] Hui Jiang, Chulun Zhou, Fandong Meng, Biao Zhang, Jie Zhou, Degen Huang, Qingqiang Wu, and Jinsong Su. 2021. Exploring Dynamic Selection of Branch Expansion Orders for Code Generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 5076–5085.
- [10] Diederik P Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*.
- [11] Wang Ling, Phil Blunsom, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, Fumin Wang, and Andrew Senior. 2016. Latent Predictor Networks for Code Generation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 599–609.
- [12] Ilya Loshchilov and Frank Hutter. 2018. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- [13] Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin Clement, Dawn Drain, Daxin Jiang, Duyu Tang, et al. 2021. CodeXGLUE: A Machine Learning Benchmark Dataset for Code Understanding and Generation. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- [14] Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective Approaches to Attention-based Neural Machine Translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. 1412–1421.
- [15] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*. 311–318.
- [16] Maxim Rabinovich, Mitchell Stern, and Dan Klein. 2017. Abstract Syntax Networks for Code Generation and Semantic Parsing. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 1139–1149.
- [17] Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural Machine Translation of Rare Words with Subword Units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 1715–1725.
- [18] Zeyu Sun, Qihao Zhu, Lili Mou, Yingfei Xiong, Ge Li, and Lu Zhang. 2019. A grammar-based structural cnn decoder for code generation. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 7055–7062.
- [19] Zeyu Sun, Qihao Zhu, Yingfei Xiong, Yican Sun, Lili Mou, and Lu Zhang. 2020. Treegen: A tree-based transformer architecture for code generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 8984–8991.
- [20] Alexey Svyatkovskiy, Shao Kun Deng, Shengyu Fu, and Neel Sundaresan. 2020. Intellico compose: Code generation using transformer. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1433–1443.
- [21] Daniel C Wang, Andrew W Appel, Jeff L Korn, and Christopher S Serra. 1997. The Zephyr abstract syntax description language. In *Proceedings of the Conference on Domain-Specific Languages on Conference on Domain-Specific Languages (DSL)*, 1997. 17–17.
- [22] Yue Wang, Weishi Wang, Shafiq Joty, and Steven CH Hoi. 2021. CodeT5: Identifier-aware Unified Pre-trained Encoder-Decoder Models for Code Understanding and Generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 8696–8708.
- [23] Binbin Xie, Jinsong Su, Yubin Ge, Xiang Li, Jianwei Cui, Junfeng Yao, and Bin Wang. 2021. Improving Tree-Structured Decoder Training for Code Generation via Mutual Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 14121–14128.
- [24] Frank F Xu, Zhengbao Jiang, Pengcheng Yin, Bogdan Vasilescu, and Graham Neubig. 2020. Incorporating External Knowledge through Pre-training for Natural Language to Code Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 6045–6052.
- [25] Pengcheng Yin and Graham Neubig. 2018. TRANX: A Transition-based Neural Abstract Syntax Parser for Semantic Parsing and Code Generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. 7–12.
- [26] Jingyi Zhang, Masao Utiyama, Eiichiro Sumita, Graham Neubig, and Satoshi Nakamura. 2018. Guiding Neural Machine Translation with Retrieved Translation Pieces. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. 1325–1335.