

Programming Language Agnostic Mining of Code and Language Pairs with Sequence Labeling Based Question Answering

Anonymous EMNLP submission

Abstract

Mining aligned natural language (NL) and programming language (PL) pairs is a critical task to NL-PL understanding. Existing methods applied specialized hand-crafted features or separately-trained models for each PL. However, they usually suffered from low transferability across multiple PLs, especially for niche PLs with less annotated data. Fortunately, a Stack Overflow answer post is essentially a sequence of text and code blocks and its global textual context can provide PL-agnostic supplementary information. In this paper, we propose a Sequence Labeling based Question Answering (SLQA) method to mine NL-PL pairs in a PL-agnostic manner. In particular, we propose to apply the BIO tagging scheme instead of the conventional binary scheme to mine the code solutions which are often composed of multiple blocks of a post. Experiments on current single-PL single-block benchmarks and a manually-labeled cross-PL multi-block benchmark demonstrate the effectiveness and transferability of SLQA. We further present a large-scale parallel NL-PL corpus named Lang2Code automatically mined with SLQA, which contains over 1.4M pairs on six PLs, for further NL-PL research.

1 Introduction

Recent years, a new suite of developer assistance tools based on natural language processing (NLP) have started to burgeon, such as code retrieval (Allamanis et al., 2015; Wei and Ch, 2015; Husain et al., 2020), code summarization (Allamanis et al., 2016; Iyer et al., 2016) and code synthesis (Locascio et al., 2016; Desai et al., 2016; Quirk et al., 2015; Yin and Neubig, 2017; Clement et al., 2020). Since data-driven methods have shown great promise on these tasks (Yin et al., 2018), high-quality parallel natural language (NL) and programming language (PL) corpora have gradually become essential resources, making NL-PL pair mining a critical task for these downstream applications.

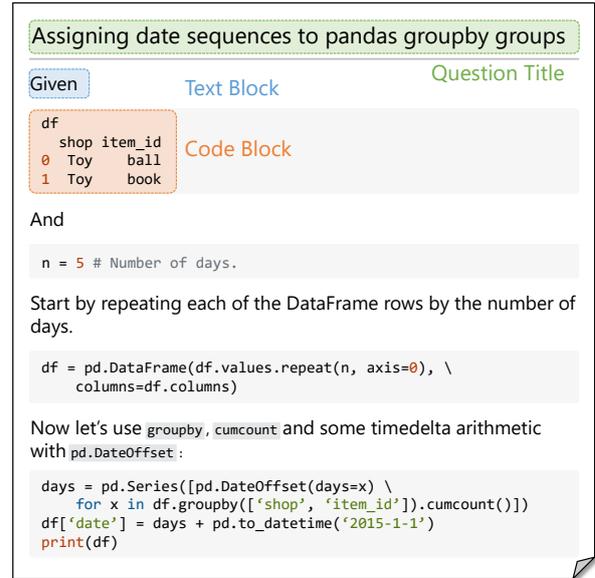


Figure 1: An example of a Stack Overflow answer post, which contains a NL title and a sequence of text and code blocks. The solution of this post consists of the last two code blocks.

Existing methods acquired the aligned NL-PL pairs through handcrafting specialized features (Iyer et al., 2016; Yin et al., 2018) or training separate models (Yao et al., 2018) for each PL. For example, Yin et al. (2018) proposed to combine PL-specific handcrafted rules with neural network models. Yao et al. (2018) proposed to incorporate the local textual context around code blocks and train separate models for Python and SQL data. However, these methods are highly dependent on large amounts of single-PL training data for each PL, making it difficult to directly migrate them to niche PLs that lack training data.

Fortunately, we observe that a Stack Overflow (SO) answer post is essentially a sequence of text and code blocks. The global textual context of the post can provide PL-agnostic supplementary information for mining correct code blocks corresponding to the post’s NL title. For example, in

Figure 1, even if we remove the content of the code blocks, the remaining PL-agnostic context “Given C_1 and C_2 , start by ... C_3 , now let’s use ... C_4 ” still can help identify C_3 and C_4 as the correct solution. This is quite similar to a question answering task (Rajpurkar et al., 2016; Yang et al., 2018), which aims to identify the correct answer of the given question from a textual context.

Inspired by the above observations, we propose a Sequence Labeling based Question Answering (SLQA) method to mine NL-PL pairs in a PL-agnostic manner. We first tokenize the text and code blocks of a post with a PL-independent tokenizer, which eliminates the characteristics of the PL itself. Then we feed the title and the converted content into a sequence tagger. The tagger assigns a label to each code-block token to indicate whether it belongs to the answer. Moreover, we observe that many solutions provided by SO posts are composed of multiple code blocks interspersed with some explanatory text, as shown in Figure 1. Therefore, we propose to apply the BIO tagging scheme (Ramshaw and Marcus, 1995) to predict multiple code blocks as answers instead of the binary scheme as in (Yao et al., 2018).

The main contribution of this paper are:

- We propose the SLQA method to mine NL-PL pairs in a PL-agnostic manner.
- We propose to apply the BIO tagging scheme instead of the previous binary scheme to mine code solutions composed of multiple blocks.
- We conduct experiments on the current single-PL single-block benchmarks and a manually-annotated cross-PL multi-block benchmark, and the results demonstrate the effectiveness and transferability of our method.
- We further present a automatically-mined parallel NL-PL corpus Lang2Code containing over 1.4M aligned pairs on six PLs.

2 Related Work

2.1 NL-PL Pair Mining

Existing methods for mining NL-PL pairs first proposed to utilize the function’s code-documentation pairs because they are naturally aligned and can be automatically mined from sources like GitHub. For example, Husain et al. (2020) proposed the CodeSearchNet corpus by collecting function-level documentations on the GitHub. Clement et al. (2020) proposed to translate documentation strings to

Python codes with transformers (Vaswani et al., 2017). Feng et al. (2020) proposed the CodeBERT model which is pre-trained with function-docstring pairs. However, language used in documentations is fundamentally different from the general natural language. A function’s documentation is often written by the same author of the code and therefore does not differ much from the code’s vocabulary (Husain et al., 2020).

To address this issue, subsequent work focused on mining NL-PL pairs from linguistically richer but unstructured code forums from sources like SO. For example, Iyer et al. (2016) proposed a neural network to generate code summaries with posts that has only one code block. Yin et al. (2018) proposed to extract structured features to explore Python and Java forums and present the automatically mined CoNaLa corpus. Yao et al. (2018) proposed a bi-view hierarchical neural network and present the StaQC dataset. However, existing methods usually required heavy PL-specific feature engineering or heavy PL-specific data annotations. This poses great challenges for directly migrating them across multiple PLs, especially for niche PLs which lack annotated data. To overcome these drawbacks, we propose a PL-agnostic way to predict multiple code blocks as a solution. We treat the code contents as natural language and tokenize all the text and code blocks with a PL-independent tokenizer, which eliminates the specific characteristics of the PLs. Then our model identifies the correct solution(s) of the question by understanding the global PL-agnostic textual context of the post.

2.2 NL Context Understanding

Recent progress of NLP techniques provided effective QA methods for comprehending query-related NL context. For example, Seo et al. (2017) proposed a BiDAF model to capture interactions between contexts and queries. Devlin et al. (2019) and Liu et al. (2019) showed improved QA performance using pre-trained language models. However, these conventional QA model usually predicted only one contiguous answer span and are not suitable for answers consisting of consecutive code blocks interspersed with text blocks. In comparison, we propose to address this issue with the sequence labeling mechanism (Lample et al., 2016). We assign a BIO tag (Ramshaw and Marcus, 1995) to each code block to indicate whether it belongs to the solution of the question title.

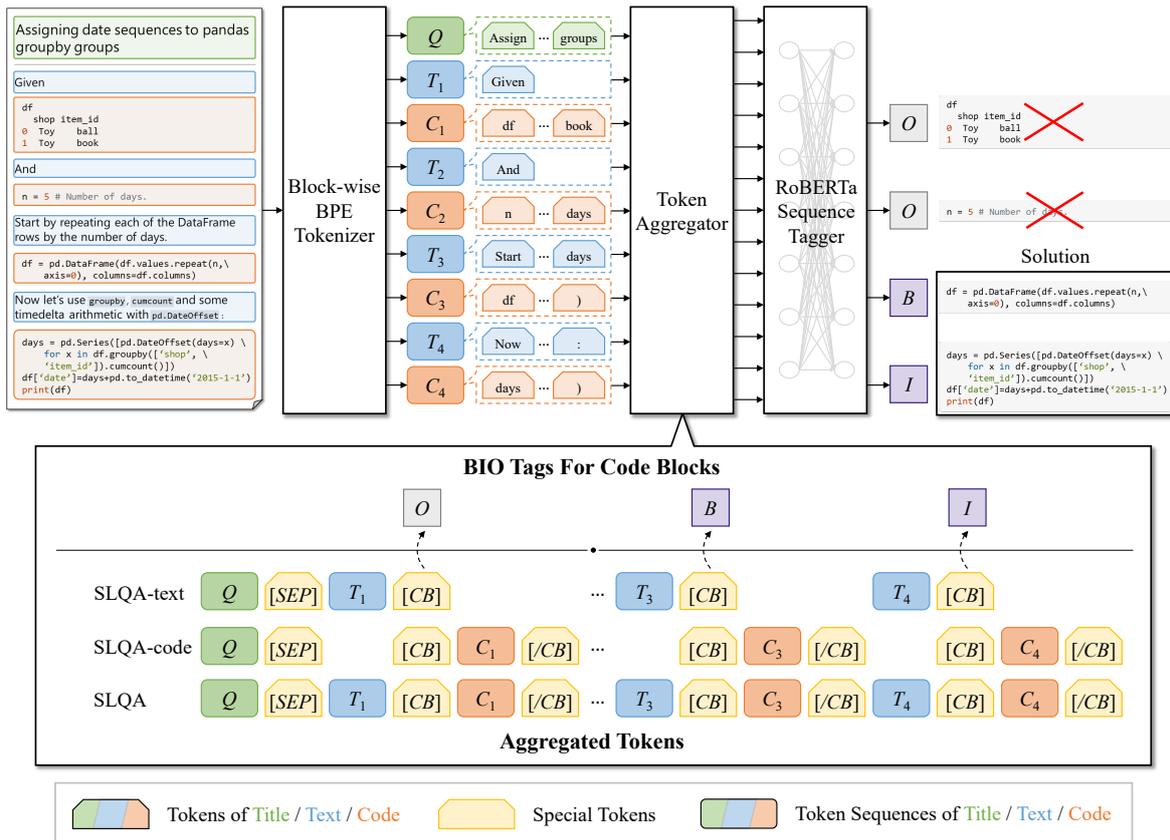


Figure 2: The overall framework of SLQA and two variants: SLQA-text and SLQA-code. The SLQA-code and SLQA models use code block contents as input. However, the codes are tokenized with the same Byte-Pair Encoding (BPE) (Gage, 1994) tokenizer as the NL context. This means that these two models are still PL-agnostic.

3 Preliminaries

In this section, we provide a task definition and describe how we annotated a cross-PL multi-block dataset for model development.

3.1 Task Overview

Given a Stack Overflow question title and its answer post accepted by the questioner, we aim at identifying all potential consecutive code blocks in the accepted answer that can solve the question. The solutions may contain one single code block or multiple code blocks interspersed with texts. Following Yao et al. (2018), we focus on "how-to-do-it" questions in SO, because their answer posts are more likely to contain complete code solutions.

3.2 Manual Annotation

We first collected questions of six PLs from Stack Overflow's raw data dump¹: Python, Java, SQL, R, Git and Bash. Then we manually selected how-to-questions by three annotators. We kept the ques-

¹<https://archive.org/details/stackexchange>

tions with accepted answers, and for each question we only kept its title and the accepted answer post following (Yao et al., 2018). Next, we annotated the standalone code blocks² in the accepted answer with BIO tags. Specifically, we use the following annotation protocol: (1) non-solution blocks are labeled as "Outside (O)"; (2) single-block solutions are labeled as "Begin (B)"; (3) for multi-block solutions, the first code block is labeled as " B " and all subsequent ones are labeled as "Inside (I)". This protocol was determined by three annotators through annotating a unified set of sample data and running five iterations of pilot study. The average Cohen's kappa of the final iteration is around 0.785. Finally, all data were exclusively divided to the three annotators for annotation.

4 Our Approach

The overall framework of our Sequence Labeling based Question Answering (SLQA) approach is

²We ignored in-line code blocks as our pilot annotation study found that standalone blocks cover most of the solutions.

shown in Figure 2. First, we introduce the intuition in Section 4.1. Then we introduce the tokenizer and the token aggregator in Section 4.2. Finally we introduce the sequence tagger in Section 4.3.

4.1 Intuition

Assuming that an answer post has a question title \mathbb{Q} and a sequence of text and code blocks: $\{\mathbb{T}_1, \mathbb{C}_1, \mathbb{T}_2, \mathbb{C}_2 \dots \mathbb{T}_N, \mathbb{C}_N\}$, where $\mathbb{T}_{1:N}$ and $\mathbb{C}_{1:N}$ denote the text and code blocks, respectively. On one hand, we observe that the global textual context of $\mathbb{T}_{1:N}$ can provide *PL-independent* information to identify the solution. For example, in Figure 1, the words "Given" and "And" of the first two text blocks $\{\mathbb{T}_1, \mathbb{T}_2\}$ strongly indicate $\{\mathbb{C}_1, \mathbb{C}_2\}$ as prerequisites. Then the keywords "Start by" and "Now" of $\{\mathbb{T}_3, \mathbb{T}_4\}$ suggest that \mathbb{C}_3 and \mathbb{C}_4 are the actual solution of the post.

On the other hand, the contents of the code blocks are intuitively informative. However, previous work utilized *PL-specific* parsers to pre-process the code (Yao et al., 2018), which is not suitable for cross-PL transfer learning. To overcome this drawback, our full model **SLQA** incorporates code contents by treating the code contents as natural language and using a unified *PL-agnostic* tokenizer (Section 4.2) to process all text and code blocks.

Surprisingly, we also discover that even without looking at the contents of all the code blocks, the remaining textual context "Given \mathbb{C}_1 and \mathbb{C}_2 , start by ... \mathbb{C}_3 , now let's use ... \mathbb{C}_4 " still can help identify the solution. Therefore, we propose a pure-text variant of SLQA, denoted as **SLQA-text**. Then another pure-code variant **SLQA-code** is naturally proposed, which is useful to investigate the effects of text and code separately, as shown in Figure 2.

4.2 Tokenizer and Token Aggregator

Given an answer post, we first use a unified tokenizer to tokenize the question title, the text blocks, and the code blocks. However, the variance of different PL's keyword sets and the flexibility of the variables' names may lead to a large but sparse vocabulary. We apply a Byte-Pair Encoding (BPE) (Gage, 1994) tokenizer to all blocks to alleviate data sparsity following (Liu et al., 2019). We denote the output token sequences of \mathbb{T}_i , \mathbb{C}_i and \mathbb{Q} as the italic symbols T_i , C_i and Q , respectively.

Then we aggregate the above tokens to construct the sequence tagger's input (Figure 2). First, we put $T_{1:N}$ and $C_{1:N}$ after Q by the order of appearance. Then, we insert two kinds of special tokens

to the whole sequence: (1) insert $[SEP]$ after Q to separate questions and contents (Devlin et al., 2019); (2) insert $[CB]$ and $[/CB]$ around each C_i to distinguish codes from texts. For SLQA-code, we remove all text tokens T_* . For SLQA-text, we remove all code tokens C_* and $[/CB]$ s, and left $[CB]$ s as placeholders of code blocks.

4.3 Sequence Taggers

Inspired by previous QA methods (Devlin et al., 2019; Liu et al., 2019), we use a pre-trained language model to encode the input token sequence, such as RoBERTa (Liu et al., 2019) in our model. We use its last hidden states $\mathbf{h}_{1:M}^S$ as the token representations of the input sequence. However, conventional QA models only predict the begin and end positions of a *contiguous* token span as the correct answer, thus failed to predict multiple code blocks interspersed with texts. For example, in Figure 2, the answer tokens C_3 and C_4 are *non-contiguous* because the token span is cut off by T_4 in the middle. To address this issue, we propose to apply the sequence labeling method for answer prediction. We predict a BIO tag for each $[CB]$ token to find multiple consecutive code blocks as the solution, as shown in Figure 2. Formally, we use a softmax classifier to predict the tag for the i -th block ($i \in \{1, 2, \dots, N\}$):

$$p(\mathbf{y}_i) = \text{softmax}(\mathbf{W}\mathbf{h}_{m_i}^S + \mathbf{b})$$

$$\text{tag}_i = \arg \max_{k \in \{B, I, O\}} p(y_i = k) \quad (1)$$

where (\mathbf{W}, \mathbf{b}) are parameters of the classifier, m_i denotes the position of the i -th $[CB]$ token. We train our model with a cross-entropy loss:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \log p(y_i = \hat{y}_i) \quad (2)$$

where \hat{y}_{m_i} denotes the gold tag of the i -th block.

Note that Yao et al. (2018) also proposed to capture textual context for this task. Our model essentially differs from theirs in that we apply a BIO tagging scheme to find multi-block solutions, whereas their binary scheme could only find single-block solutions. Moreover, we model the post content as a whole sequence and capture its global textual context, while their methods only incorporated the local textual context around the code blocks. We show that the global textual context provides more information than the local context in Section 5.4.

5 Experiments

5.1 Datasets

We evaluate our methods on two datasets, which are referred to as StaQC-human and Lang2Code-human respectively in the following sections.

5.1.1 StaQC-human

StaQC (Yao et al., 2018) consists of 4,883 Python and 3,637 SQL manually annotated question-code (QC) pairs. Each code block is annotated with 0 or 1 to indicate whether it solves the corresponding question. This means that StaQC discards all multi-block joint solutions. The original training, validation and test sets are split at the code block level and therefore not applicable to our post-level approach. Moreover, two consecutive code blocks from the same post may be placed into different subsets. Since they share the same text block between them, this placement may cause potential data leakage. Therefore, we reassemble the code blocks into posts and randomly split the fully-annotated posts into new training, validation and test sets. Then we transform the $\{1, 0\}$ labels into $\{B, O\}$ tags. We denote the new single-block dataset as StaQC-human. Table 1 shows the statistics of the single-block StaQC-human dataset.

5.1.2 Lang2Code-human

We construct the Lang2Code-human dataset with our manually-annotated data (Section 3.2). We focused on Python as the primary PL. The Python posts are used as the training, validation sets and part of the test set. They contain both randomly-sampled and hits-weighted sampled data to jointly represent the distributions of all SO posts and the popular ones. Posts of other five PLs (Java, SQL, R, Git, Bash) are randomly sampled and used as the

Subset	Lang.	#Post	#Block	
			Single	Non.
Train	Python	1,139	1,354	1,592
Valid	Python	142	181	197
Test	Python	143	164	204
Train	SQL	1,056	1,632	1,016
Valid	SQL	132	201	124
Test	SQL	133	190	135

Table 1: Statistics of the StaQC-human dataset. Non. stands for non-solution blocks.

Subset	Lang.	#Post	#Block		
			Single	Multi.	Non.
Train	Python	1,058	1,328	423	1,131
Valid	Python	132	185	58	140
	Python	133	178	44	149
	Java	53	47	68	23
	SQL	70	91	33	60
	R	80	88	44	84
	Git	75	63	107	41
Test	Bash	69	87	63	55

Table 2: Statistics of the Lang2Code-human dataset. *Multi.* stands for code blocks that form multi-block solutions, and *Non.* stands for non-solution blocks.

test set for cross-PL performance benchmarking. The statistics of the multi-block Lang2Code-human dataset is shown in Table 2.

5.2 Experimental Settings

We tune the hyper-parameters on the validation sets with earlystopping patience of 5. We load the pre-trained RoBERTa-Large parameters³ and fine-tune them during training. We use the Adam (Kingma and Ba, 2014) optimizer to train our model with the learning rate of 10^{-5} and 500 warmup steps. We set the weight decay to 0.01. We report the precision, recall and F_1 scores of the predict solutions on both datasets. We also report the accuracy on StaQC-human because of its binary labels. We train our model for five times with different random seeds and report the average scores.

5.3 Baselines

We compare our model with the following baselines: (1) **Select-All** and **Select-First** are two commonly used heuristics (Yao et al., 2018). Select-All treats all code blocks in a post as standalone solutions, i.e., all code blocks are labelled as "B". Select-First only chooses the first code block as the standalone solution, which is equivalent to labelling the first code block as "B" and the rest as "O". (2) **BiV-HNN** (Yao et al., 2018) learns the semantic representation of a code block from token level to block level and incorporates the local textual context around each code block. It has two main variants, **Text-HNN** and **Code-HNN**, which discard code content and textual context, respectively. Although StaQC contains data of two PLs,

³<https://huggingface.co/roberta-large>

Method	Python					SQL				
	Train	Prec.	Rec.	F_1	Acc.	Train	Prec.	Rec.	F_1	Acc.
Select-First	Python	64.3	56.1	60.0	66.6	SQL	73.7	51.3	60.5	60.6
Select-All	Python	44.6	100	61.7	44.6	SQL	58.8	100	74.0	58.8
Text-HNN	Python	71.7	79.5	75.4	76.8	SQL	76.5	86.6	81.2	76.5
Code-HNN	Python	71.7	77.8	74.6	76.4	SQL	75.3	88.8	81.5	76.3
BiV-HNN	Python	76.2	83.8	79.6	81.0	SQL	83.2	96.0	89.1	86.1
BiV-RoBERTa	Python	75.8	76.1	74.7	77.7	SQL	82.9	91.1	86.6	83.4
SLQA-text	Python	78.5	86.0	<u>82.0</u>	<u>83.0</u>	SQL	86.3	94.0	<u>90.0</u>	<u>87.7</u>
SLQA-code	Python	72.8	87.0	79.1	79.1	SQL	84.2	89.0	86.5	83.5
SLQA	Python	85.3	<u>90.2</u>	87.6	88.4	SQL	85.2	<u>97.7</u>	91.0	88.4
SLQA-text [†]	SQL	74.6	90.1	81.6	81.9	Python	<u>88.0</u>	82.8	85.2	83.1
SLQA-code [†]	SQL	<u>84.5</u>	24.3	37.3	63.6	Python	77.3	80.1	78.3	74.0
SLQA [†]	SQL	79.5	82.3	80.9	82.3	Python	90.4	88.3	89.2	87.2

Table 3: Performance of SLQA and previous method on StaQC-human’s Python and SQL test sets. The best scores are in bold and the second-best scores are underlined. † denotes cross-PL performance.

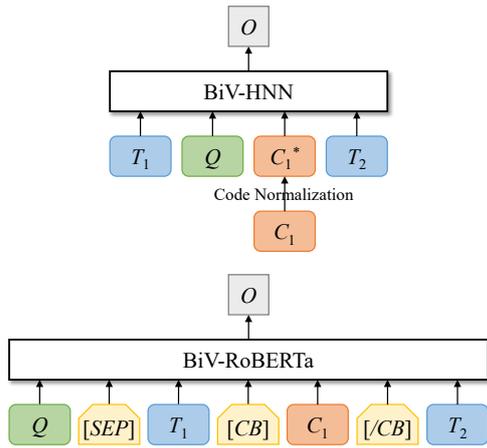


Figure 3: The overall framework of BiV-HNN and BiV-RoBERTa. BiV-HNN is a GRU based model that takes normalized code as input, whereas BiV-RoBERTa is a transformer based model that takes raw code as input.

these models are not available for cross-PL experiments because they used PL-specific vocabularies. Since StaQC-human is re-split at the post level, we re-train BiV-HNN and its variants on StaQC-human with the official implementation and report the post-level scores. (3) **BiV-RoBERTa** is similar to BiV-HNN, but uses the RoBERTa instead of the HNN to capture the code contents and local textual context, as shown in Figure 3.

5.4 StaQC-human Results

We conduct experiments on StaQC-human Python dataset and report the results in Table 3, from which we have the following observations:

Single-PL: (1) SLQA significantly outperforms previous Select-First, Select-All and BiV-HNN methods. It indicates that our model can effectively capture global textual context and code contents and improve the performance. (2) SLQA outperforms the BiV-RoBERTa model. The two models use the same RoBERTa encoder, the only difference is that SLQA uses global contextual information while the other uses local contextual information. It indicates the effectiveness of the global textual context for NL-PL understanding. Additionally, BiV-RoBERTa performs worse than BiV-HNN. It shows that the PL-specific code normalization (Figure 3) provides major contribution to the performance. SLQA performs even better than BiV-HNN without code normalization, which indicates that the improvements come primarily from the global textual context rather than the pre-trained RoBERTa. (3) SLQA-text also outperforms the BiV-HNN model. Since SLQA-text has no access to the code contents but still brings improvements, this result further demonstrates the effectiveness of our method to capture global textual information.

Cross-PL: (1) SLQA-text achieves comparable results to the single-PL BiV-HNN model. It indicates that the global textual context provides PL-agnostic supplementary information and allows for transfer learning across different PLs. (2) SLQA-code performs well on Single-PL task but poorly on Cross-PL task, even worse than the heuristic baselines on the Python test set. It validates the importance of textual information to cross-PL transfer

Method	F_1					
	Python	Java	SQL	R	Git	Bash
Select-First	49.8	37.5	40.2	45.2	30.2	41.3
Select-All	63.4	44.1	63.7	54.5	41.1	55.9
SLQA-text	69.7	<u>53.8</u>	<u>70.0</u>	64.0	<u>55.4</u>	<u>65.8</u>
SLQA-code	<u>70.6</u>	50.0	59.8	<u>64.7</u>	48.0	61.5
SLQA	75.7	62.9	76.1	66.5	58.7	70.1

Table 4: Performance of SLQA and heuristics on Lang2Code-human test set, including all the six PLs. The best scores are in bold and the second-best scores are underlined. Our models are trained on Python training set and evaluated on the 6 test sets to examine the transfer learning ability.

learning. (3) Cross-PL SLQA outperforms single-PL BiV-HNN and achieves comparable results to single-PL SLQA-text. It indicates that our model can effectively incorporate the code contents and textual context in a PL-agnostic manner and improve the cross-PL transfer learning performance.

5.5 Lang2Code-human Results

We present the experimental results on the cross-PL multi-block Lang2Code dataset in Table 4. We can observe that: (1) The SLQA scores on the Python test set are significantly lower than StaQC-human’s single-PL scores. It demonstrate that predicting multi-block solutions is a more challenging task than predicting single-block solutions. (2) SLQA-code performs poorly on cross-PL learning except for Python-to-R learning. It indicates that only focusing code contents is not effective enough for cross-PL learning because the language style varies greatly between different PLs, unless the two PLs have similar styles like Python and R. (3) SLQA-text significantly outperforms heuristic methods. It indicates that the global textual context can provide critical PL-agnostic information to identify the correct solution from an SO post. Our model effectively captures the global textual context and improve the performance of cross-PL transfer learning. (4) The full SLQA model significantly outperforms both the variants and the heuristics. It indicates that SLQA effectively captures global textual context and code contents of a post and allows for reliable transfer learning across multiple PLs.

6 Lang2Code: An automatically mined NL-PL pairs dataset

In this section, we describe how to automatically mine NL-PL pairs with SLQA and obtain Lang2Code, a large-scale parallel NL-PL dataset.

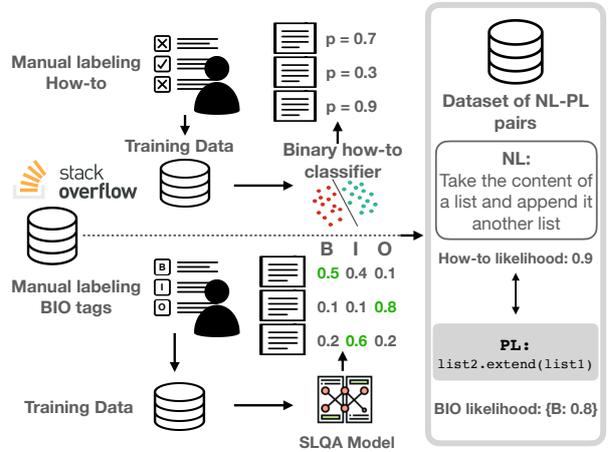


Figure 4: The overall pipeline of automatically constructing the Lang2Code dataset.

The overall pipeline is shown in Figure 4.

In Section 3.2, we obtained 2,351 manually annotated tags from Lang2Code for the six languages, concomitant with these tags are also binary labels indicating whether the question is a How-to question. Using this data, we fine-tuned BERT⁴ to train a binary classifier, obtaining 87.4% precision, 89.0% recall and 89.9% f1 score.

Next we used this classifier with a threshold of 0.5 to filter Stack Overflow posts pertaining to the six PLs we annotated in Section 3.2. We filtered 1,000,888 single-block answer posts and 799,104 multi-block answer posts and kept 540,352 and 676,128 how-to questions respectively.

After that, for single-block answer posts, we follow previous work (Iyer et al., 2016; Yin et al., 2018; Yao et al., 2018) to treat the only code block as the solution to the question title. And thus we can have a NL-PL pair out of each single-code

⁴<https://huggingface.co/bert-base-cased>

Lang.	#Pairs	#Solutions		#Blocks	
		Single	Multi.	Single	Multi.
Python	572,587	537,768	34,819	537,768	81,574
Java	322,731	295,920	26,811	295,920	62,487
SQL	234,939	226,548	8,391	226,548	19,040
R	195,685	181,862	13,823	181,862	33,229
Git	27,470	24,164	3,306	24,164	8,549
Bash	45,099	43,104	1,995	43,104	4,801
Total	1,398,511	1,309,366	89,145	1,309,366	209,680

Table 5: The statistics of Lang2Code. #Solution Multi. refers to the number of NL-PL pairs that PL consists of multiple code blocks. #Blocks Multi. refers to the total number of code blocks combined in those NL-PL pairs that consist of multiple code blocks.

URL: <https://stackoverflow.com/questions/59792278>

How do I list all local tracking branches not pushed in git? H:0.94

Solution 1

```
git fetch
```

B:0.63

```
git branch -vv | grep origin/
```

I:0.72

```
*master 083378cc35 [origin/master]...
xyz      8dbdf339cd ...
```

O:0.88

Solution 2

```
git rev-list --count --left-right ...
```

B:0.47

Figure 5: An illustration of the Lang2Code Corpus. From a SO post, two solutions corresponding to the question title were obtained, resulting in two NL-PL pairs. In the corpus, both the how-to likelihood and the BIO likelihoods are included.

answer posts. And for posts contain multiple code blocks, we applied SLQA trained on Lang2Code-human to the question answer pair and assigned BIO tags to every code blocks.

As discussed in Section 3.2, using BIO tags we can segregate multiple code-block solutions as the ones with "B" followed by one or more "I" tags, whereas single code-block solutions are ones with just a "B" tag followed by "O" tags or is the last code block. We pair each code-block solution with the question title to construct an NL-PL pair. For example, in figure 5, We use Solution 1 and Solu-

tion 2 to pair with the question title separately to obtain two NL-PL pairs.

With this approach, we have mined 858,159 NL-PL pairs from multi-code answer posts, together with the NL-PL pairs obtained from single-code answer posts, making a total of 1,398,511 NL-PL pairs in total. To the best of our knowledge, this is the largest NL-PL corpus mined from Stack Overflow so far, and the first corpus containing code solutions that are composed of multiple code blocks. We name this corpus Lang2Code, and display its statistics in table 5. We will in the future publish the Lang2Code corpus and open source all related code for future research.

7 Conclusion

In this paper, we propose a Sequence Labeling based Question Answering (SLQA) approach to construct NL-PL pairs in a PL-agnostic way. We propose a model that can produce multiple code blocks as solutions of a post's question, which is achieved by using BIO sequence tagging. We also propose to incorporate the global textual context as PL-independent supplementary information. To validate the capacity of our method, we manually annotate a challenging cross-PL multi-block dataset, named Lang2Code-human. Substantial experiments on the single-PL single-block StaQC-human and our Lang2Code-human benchmarks demonstrate the effectiveness and cross-PL transferability of our method. Finally, we present Lang2Code, the largest-to-date NL-PL corpus to the best of our knowledge, containing over 1.4 million pairs spanning six PLs.

References

- Miltiadis Allamanis, Hao Peng, and Charles Sutton. 2016. [A convolutional attention network for extreme summarization of source code](#).
- Miltos Allamanis, Daniel Tarlow, Andrew Gordon, and Yi Wei. 2015. [Bimodal modelling of source code and natural language](#). In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2123–2132, Lille, France. PMLR.
- Colin B. Clement, Dawn Drain, Jonathan Timcheck, Alexey Svyatkovskiy, and Neel Sundaresan. 2020. [Pytm5: multi-mode translation of natural language and python code with transformers](#).
- Aditya Desai, Sumit Gulwani, Vineet Hingorani, Nidhi Jain, Amey Karkare, Mark Marron, and Subhajit Roy. 2016. [Program synthesis using natural language](#). In *Proceedings of the 38th International Conference on Software Engineering*, pages 345–356.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). In *NAACL-HLT*, pages 4171–4186.
- Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, and Ming Zhou. 2020. [Codebert: A pre-trained model for programming and natural languages](#).
- Philip Gage. 1994. [A new algorithm for data compression](#). *C Users J.*, 12(2):23–38.
- Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. 2020. [Code-searchnet challenge: Evaluating the state of semantic code search](#).
- Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, and Luke Zettlemoyer. 2016. [Summarizing source code using a neural attention model](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2073–2083, Berlin, Germany. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *arXiv preprint arXiv:1412.6980*.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. [Neural architectures for named entity recognition](#). *arXiv preprint arXiv:1603.01360*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Nicholas Locascio, Karthik Narasimhan, Eduardo DeLeon, Nate Kushman, and Regina Barzilay. 2016. [Neural generation of regular expressions from natural language with minimal domain knowledge](#).
- Chris Quirk, Raymond Mooney, and Michel Galley. 2015. [Language to code: Learning semantic parsers for if-this-then-that recipes](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 878–888, Beijing, China. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100,000+ questions for machine comprehension of text](#).
- Lance Ramshaw and Mitch Marcus. 1995. [Text chunking using transformation-based learning](#). In *Third Workshop on Very Large Corpora*.
- Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. [Bidirectional attention flow for machine comprehension](#).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *NIPS*, pages 6000–6010.
- Yi Wei and Nirupama Ch. 2015. [Building bing developer assistant](#).
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [Hotpotqa: A dataset for diverse, explainable multi-hop question answering](#).
- Ziyu Yao, Daniel S. Weld, Wei-Peng Chen, and Huan Sun. 2018. [Staqa. Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW '18](#).
- Pengcheng Yin, Bowen Deng, Edgar Chen, Bogdan Vasilescu, and Graham Neubig. 2018. [Learning to mine aligned code and natural language pairs from stack overflow](#).
- Pengcheng Yin and Graham Neubig. 2017. [A syntactic neural model for general-purpose code generation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 440–450, Vancouver, Canada. Association for Computational Linguistics.