

CODE-MVP: Learning to Represent Source Code from Multiple Views with Contrastive Pre-Training

Xin Wang¹[◇] Yasheng Wang² Yao Wan³ Jiawei Wang⁴

Pingyi Zhou² Li Li⁴ Hao Wu⁵ Jin Liu¹[✉]

¹School of Computer Science, Wuhan University, China ²Huawei Noah’s Ark Lab

³School of Computer Sci. & Tech., Huazhong University of Science and Technology, China

⁴Faculty of Information Technology, Monash University, Australia

⁵School of Information Science and Engineering, Yunnan University, China

{xinwang0920, jinliu}@whu.edu.cn

Abstract

Recent years have witnessed increasing interest in code representation learning, which aims to represent the semantics of source code into distributed vectors. Currently, various works have been proposed to represent the complex semantics of source code from different views, including plain text, Abstract Syntax Tree (AST), and several kinds of code graphs (e.g., Control/Data Flow Graph). However, most of them only consider a single view of source code independently, ignoring the correspondences among different views. In this paper, we propose to integrate different views with the natural-language description of source code into a unified framework with **Multi-View** contrastive **Pre-training**, and name our model as **CODE-MVP**. Specifically, we first extract multiple code views using compiler tools, and learn the complementary information among them under a contrastive learning framework. Inspired by the type checking in compilation, we also design a fine-grained type inference objective in the pre-training. Experiments on three downstream tasks over five datasets demonstrate the superiority of **CODE-MVP** when compared with several state-of-the-art baselines. For example, we achieve 2.4/2.3/1.1 gain in terms of MRR/MAP/Accuracy metrics on natural language code retrieval, code similarity, and code defect detection tasks, respectively.

1 Introduction

Code intelligence that utilizes machine learning techniques to promote the productivity of software developers, has attracted increasing interest in both communities of software engineering and artificial intelligence (Lu et al., 2021; Feng et al., 2020; Wang et al., 2022; Wan et al., 2022a; Wu et al., 2021). To achieve code intelligence, one fundamental task is code representation learning (also

[◇] Work conducted during an internship at Huawei Noah’s Ark Lab.

[✉] Corresponding author.

Models	Tokens	AST	Graph	PT
CodeBERT (Feng et al., 2020)	✓	✗	✗	✗
GraphCodeBERT (Guo et al., 2021)	✓	✗	✓	✗
SynCoBERT (Wang et al., 2021)	✓	✓	✗	✗
CodeGPT (Lu et al., 2021)	✓	✗	✗	✗
PLBART (Ahmad et al., 2021)	✓	✗	✗	✗
TreeBERT (Jiang et al., 2021)	✓	✓	✗	✗
ContraCode (Phan et al., 2021)	✓	✗	✗	✓
CoTexT (Phan et al., 2021)	✓	✗	✗	✗
CodeT5 (Wang et al., 2021b)	✓	✗	✗	✗
CODE-MVP (Our work)	✓	✓	✓	✓

Table 1: Comparison with current pre-trained code models. PT: Program Transformation.

known as code embedding), which aims to preserve the semantics of source code in distributed vectors (Alon et al., 2019). It can support various downstream tasks about code intelligence, including code defect detection (Omri and Sinz, 2020; Zhao et al., 2021b,a), code summarization (Wan et al., 2018), code retrieval (Wan et al., 2019), and code clone detection (White et al., 2016).

Current approaches to code representation borrow ideas from the successful deep learning methods in natural language processing, mainly attributed to the *naturalness hypothesis* in source code (Allamanis et al., 2018). From our investigation, existing approaches mainly represent the source code from different views of code, including code token in plain text (Iyer et al., 2016), Abstract Syntax Tree (AST) (Bui et al., 2021a), and Control/Data Flow Graphs (CFGs/DFGs) of code (Cummins et al., 2020; Wang and Su, 2020). Recently, many attempts have been made to pre-train a masked language model for source code, such as CodeBERT (Feng et al., 2020), GraphCodeBERT (Guo et al., 2021), SynCoBERT (Wang et al., 2021), CodeGPT (Lu et al., 2021), PLBART (Ahmad et al., 2021), CoTexT (Phan et al., 2021), and CodeT5 (Wang et al., 2021b). Table 1 shows the contribution of our work when compared with current pre-trained language models for source code.

Despite much progress in code representation learning, most of them only consider a single view of source code independently, ignoring the consistency among different views (Feng et al., 2020; Lu et al., 2021; Ahmad et al., 2021; Wang et al., 2021b). Usually, a program, accompanied by a corresponding natural-language comment (NL), can be parsed into multiple views, e.g., the source code tokens, AST, and CFG. We argue that these different views contain complementary semantics of the program. For example, the source code tokens (e.g., method name identifiers) and natural-language comments always reveal the *lexical* semantics of code, while the intermediate structures of code (e.g., AST and CFG) always reveal the *syntactic* and *executive* information of code. In addition, a program can also be transformed (or rewritten) into different variants that have equivalent functionality. We think that different variants of the same program reveal the *functional* information of code. That is, those different program variants with the same functionality are expected to represent the same semantics.

Inspired by the aforementioned insights, this paper proposes a novel CODE-MVP for code representation, which aims to integrate multiple views of the code into a unified framework with multi-view contrastive pre-training. Concretely, we first extract multiple views of code using several compiler tools, and learn the complementary information among them under a multi-view contrastive learning framework. Meanwhile, inspired by the type checking in compilation process, we also introduce fine-grained type inference as an auxiliary task in the pre-training process to encourage the model to learn more fine-grained type information.

To summarize, the contributions of this paper are two-fold: (1) We are the first to represent source code from multiple views, including the code tokens, AST, CFG, and various program equivalents, under a unified multi-view contrastive pre-training framework. Meanwhile, we also introduce an auxiliary task of inferring type annotations for variables. (2) We extensively evaluate CODE-MVP on three program comprehension tasks. Experimental results demonstrate the superiority of CODE-MVP when compared with several state-of-the-art baselines. Specifically, CODE-MVP achieves 2.4/2.3/1.1 gain on MRR/MAP/Accuracy metrics in natural language code retrieval, code similarity, and code defect detection tasks, respectively.

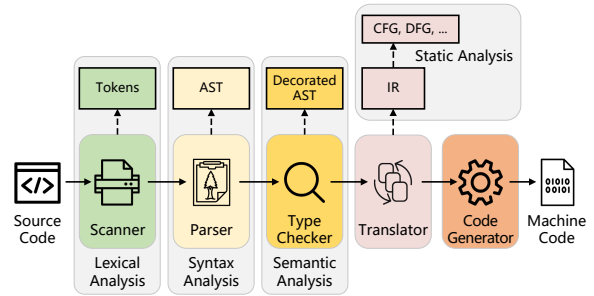


Figure 1: An example of converting a program from source code into machine code in compilation process.

2 Multiple Views of Code

We borrow ideas from the way that computers process the source code in compilation, where a program would be converted into multiple views. Figure 1 shows the process of converting a program from source code to machine code. During this process, the compiler would automatically utilize some program analysis techniques to verify the correctness of source code, including lexical, syntax, and semantic analyses. In the lexical analysis, a program is treated as a sequence of tokens and checked for spelling problems. In the syntax analysis, syntactic rules of programs are defined by the context-free grammar (Javed et al., 2004). Then the program could be parsed as an AST, based on which many program transformation heuristics can be applied to rewrite the program while maintaining the same desired functionality. In the semantic analysis, semantic rules of the program are defined by the attribute grammar (Paakki, 1995). Then the compiler could check the types of code tokens, and a decorated AST could be obtained. After the three stages above, a translator will convert the source code to its Intermediate Representation (IR), which is then considered as the basis for building Control/Data Flow Graphs (CFGs/DFGs) for further optimizations in the static analysis. Finally, the IR of the source code should be converted into machine code to execute through a code generator. Next, we introduce how we extract different views of the source code. Figure 2 illustrates multiple views of source code with an example.

Abstract Syntax Tree (AST). An AST, which is composed of leaf nodes, non-leaf nodes and edges between them, contains rich syntactic structural information of source code. In the AST, an assignment statement $y = 0$ can be represented by a non-leaf node `assignment` that points to three




Figure 2: Multiple views of source code.

leaf nodes (`0`, `y`, and `=`). In this paper, we parse a snippet of source code into an AST using a standard compiler tool `tree-sitter`.¹ To feed an AST into our model, we apply depth-first traversal to convert it into a sequence of AST tokens (Kim et al., 2021).

Control Flow Graph (CFG). CFG, which represents the execution semantics of the program in the form of a graph, is one intermediate representation of programs. A CFG consists of basic blocks and directed edges between them, where each directed edge reflects the execution order of the two basic blocks in the program. We can easily traverse the CFG along directed edges to parse it into a token sequence, which reveals the execution semantics of the program. In this paper, we use a static analyzer `Scalpel`² (Li et al., 2022) to construct the CFGs for Python code snippets.

Program Transformation (PT). The program transformation operations aim to produce multiple variants for a given program that satisfy the same desired functionality (Rabin et al., 2020). These different variants of a program can help the model capture functional semantics. In this work, we employ the following program transformation heuristics on ASTs and rewrite one program into another equivalent variant.

- **Function and Variable Renaming.** We randomly take new names from a set of candidates, such as `VAR_i`, `FUNC_i`, to rename the names of variables and functions in a program. This heuristic will not change the AST structure of the program, except for the textual appearance of variable and function names in the AST.

- **Loop Exchange.** The `for` and `while` loops represent the same functionality in a program. We traverse the AST to identify the `for` and `while` loop nodes, and replace `for` loops with `while` loops or vice versa. We also modify the initialization, condition and afterthought simultaneously.
- **Dead Code Insertion.** We first traverse the AST to identify several basic blocks (Mendis et al., 2019), and then randomly select a basic block and insert dead code snippets into it. Note that the dead code snippets are predefined and selected from a set of candidates.

3 CODE-MVP

3.1 Tasks and Notations

We define the set of program samples in multiple views (i.e. NL, PL, AST, CFG, PT) as $S = \{S^1, \dots, S^m\}$, where m represents the number of views, $s_i^a \in S^a$ represents a program in the view of a . Given a program, the PL view denotes its textual appearance, the NL view denotes its corresponding natural-language comment, and the PT denotes the variants of this program based on program transformation. The AST and CFG are extracted from a program using several compiler tools. CODE-MVP adopts two forms of input, i.e., single-view input $x_i^a = \{\langle \text{CLS} \rangle, s_i^a\}$ and dual-view input $x_i^{ab} = \{\langle \text{CLS} \rangle, s_i^a, \langle \text{SEP} \rangle, s_i^b\}$, where a and b denote two different views of the program. Following (Devlin et al., 2019), a special token `<CLS>` is appended at the beginning of each input sequence, and `<SEP>` is used to concatenate two sequences. Subsequently, the representation of `<CLS>` is used to represent the entire sequence, and `<SEP>` is used to split two views of sub-sequences. Given a set of programs with their corresponding multiple

¹<https://github.com/tree-sitter/tree-sitter>

²<https://github.com/SMAT-Lab/Scalpel>




Figure 3: An illustration of our proposed multi-view contrastive pre-training framework.

views, we aim to learn the code representation by utilizing the mutual information existing in different views. Our intuition is to learn complementary information from multiple views of code by pulling the code under different views together and pushing the dissimilar ones apart.

3.2 Framework Overview

Figure 3 shows a simple example of our multi-view contrastive pre-training framework. Given a program s_i , we use the same program to construct a pair of positive samples ($x_i^a = \{\langle \text{CLS} \rangle, s_i^a\}$ vs $x_i^b = \{\langle \text{CLS} \rangle, s_i^b\}$) in the form of views a and b , as described above. We take x_i^a and x_i^b as the input of CODE-MVP respectively. The last hidden representations of $\langle \text{CLS} \rangle$ tokens in the two inputs can be formulated as $h_i^a = \text{CODE-MVP}(x_i^a)$ and $h_i^b = \text{CODE-MVP}(x_i^b)$. We utilize a projection head (a two-layer MLP) to map hidden representations to a space, i.e., $v_i^a = f(h_i^a)$, $v_i^b = f(h_i^b)$. Then the multi-view contrastive objective can be performed. During the pre-training process, we also design other two pre-training tasks, i.e., fine-grained type inference (FGTI) task and multi-view masked language modeling (MMLM).

3.3 Multi-View Contrastive Learning

We train CODE-MVP with *paired data* and *unpaired data*. Paired data refers to those program samples with paired NL, while unpaired data stands for those isolated program samples without paired NL. Next, we explain how we construct positive

and negative samples for these two cases.

Multi-View Positive Sampling. We design *Single-View* (for paired and unpaired data) and *Dual-View* (for paired data only, which needs the NL) methods to construct multi-view positive samples for the MVCL objective:

- **Single-View.** To bridge the gap between different views of a same program, we consider the view of a program x_i^a as a positive sample w.r.t another view x_i^b . That is, ($x_i^a = \{\langle \text{CLS} \rangle, s_i^a\}$ vs $x_i^b = \{\langle \text{CLS} \rangle, s_i^b\}$) forms an inter-view positive pair, since x_i^a and x_i^b are two different views of a same program x_i .
- **Dual-View.** There are a total of C_m^2 combinations for two views of a same program. For efficiency, we focus on the features of the program itself, and propose the *NL*-conditional dual-view contrastive pre-training strategy, freezing the position of NL. Concretely, we construct a *NL*-conditional inter-view positive pair by replacing the second view in the input $\{\langle \text{CLS} \rangle, s_i^{\text{NL}}, \langle \text{SEP} \rangle, s_i^a\}$ to be $\{\langle \text{CLS} \rangle, s_i^{\text{NL}}, \langle \text{SEP} \rangle, s_i^b\}$, where $\forall a, b \neq \text{NL}$.

It is worth mentioning that there are many combinations to construct positive pairs. Some combinations are not considered in this work, such as the AST vs PT of the same program, and the CFG vs PT of the same program. Simultaneously, for training efficiency and downstream applications, we comprehensively consider eight combinations. They are (1) single-view: (NL vs PL), (NL vs PT), (PL vs AST), (PL vs CFG), and (PL vs PT); and (2) dual-view: (NL-PL vs NL-AST), (NL-PL vs NL-CFG), and (NL-PL vs NL-PT).

Multi-View Negative Sampling. Since the processes of unpaired data and paired data are similar, here we take the unpaired data as an example. We leverage *in mini-batch* and *cross mini-batch* sampling strategies (Chen et al., 2020) to construct intra-view and inter-view negative samples, respectively. Given a mini-batch of training data $b_1 = [x_1^a, \dots, x_n^a]$ in the view of a with size n , we can easily get another positive mini-batch data $b_2 = [x_1^b, \dots, x_n^b]$ in the view of b , where (x_i^a vs x_i^b) denotes an inter-view positive pair. For x_i^a , the intra-view negative samples are $\{x_j^a\}, \forall i \neq j$, and the inter-view negative samples are $\{x_j^b\}, \forall i \neq j$. Finally, for each x_i , we can get a set of $2n - 2$ negative samples.




Figure 4: Pre-training with fine-grained type inference and multi-view masked language modeling.

For an input x_i^a with representation v_i^a under the view of a , it has one positive sample x_i^b with representation v_i^b under the view of b . It also has a negative sample set $\mathbf{V}^- = \{v_1^-, \dots, v_{2n-2}^-\}$ with size $2n-2$, which consists of two types of negative sample subsets, e.g., intra-view negative sample set \mathbf{V}_1^- with size $n-1$, where $v_j^a \in \mathbf{V}_1^-, \forall j \neq i$, and the inter-view negative sample set \mathbf{V}_2^- with size $n-1$, where $v_j^b \in \mathbf{V}_2^-, \forall j \neq i$. We define the similarity of a pair of samples as the dot product of their representations. Then the loss function for a positive pair (x_i^a, x_i^b) can be defined as:

$$l(x_i^a, x_i^b) = -\ln \frac{\exp(v_i^a \cdot v_i^b)}{\exp(v_i^a \cdot v_i^b) + \sum_{k=1}^{2n-2} \exp(v_i^a \cdot v_k^-)}. \quad (1)$$

We calculate the loss for the same pair twice with order switched, i.e., (x_i^a, x_i^b) is changed to (x_i^b, x_i^a) as the dot product with negative samples for x_i^a and x_i^b are different. Overall, the MVCL loss function is defined as follows:

$$\mathcal{L}_{\text{MVCL}} = -\frac{1}{|\mathcal{N}|} \sum_i^{|\mathcal{N}|} \left[l(x_i^a, x_i^b) + l(x_i^b, x_i^a) \right], \quad (2)$$

where \mathcal{N} denotes the set of all program samples covering all different views.

3.4 Pre-Training with Type Inference

Figure 4 shows the other two pre-training tasks, including fine-grained type inference and multi-view masked language modeling.

Fine-Grained Type Inference. Several previous works (Wang et al., 2021; Wang et al., 2021b) have proven the importance of symbolic properties in programming languages. Two concurrent works, SynCoBERT (Wang et al., 2021) and CodeT5 (Wang et al., 2021b) let the model divide the code token types into `identifier` or

`non-identifier`. Inspired by the type checking in compilation process, we propose a fine-grained type inference (FGTI) objective to capture the fine-grained type information of variables (Li et al., 2022; An et al., 2011). First, we parse all source codes into ASTs. Then, we traverse the AST and use the type checker to obtain fine-grained identifier types. We employ BPE tokenizer (Sennrich et al., 2016) to tokenize tokens and let sub-tokens inherit the type information of the token. Finally, we define the loss function as follows:

$$\mathcal{L}_{\text{FGTI}} = -\frac{1}{|\mathcal{Z}|} \sum_i^{|\mathcal{Z}|} \sum_j^{|\mathcal{T}|} Y_{ij} \log P_{ij}, \quad (3)$$

where \mathcal{Z} denotes the set of all tokens that need to inference types, \mathcal{T} represents the set of all types contained in the pre-training corpus, Y_{ij} denotes the label of token i in type j , and P_{ij} denotes the predicted probability of token i in type j .

Multi-View Masked Language Modeling. In addition to the multi-view contrastive learning objective and fine-grained type inference objective, we also extend the Masked Language Modeling (MLM) to the multi-view program corpus, named MMLM. Given a data point x , we randomly select 15% of tokens in x and replace them with a special token `<MASK>`, following the same settings in (Devlin et al., 2019). The MMLM objective aims to predict original tokens which are masked out. We calculate the MMLM loss as follows:

$$\mathcal{L}_{\text{MMLM}} = -\frac{1}{|\mathcal{M}|} \sum_i^{|\mathcal{M}|} \sum_j^{|\mathcal{V}|} Y_{ij} \log P_{ij}, \quad (4)$$

where \mathcal{M} denotes the set of masked tokens, \mathcal{V} represents the vocabulary, Y_{ij} denotes the label of the masked token i in class j , and P_{ij} denotes the predicted probability of token i in class j .

3.5 Overall Training Objective

The overall loss function in CODE-MVP is the integration of several components we have defined before.

$$\mathcal{L} = \mathcal{L}_{\text{MVCL}} + \mathcal{L}_{\text{FGTI}} + \mathcal{L}_{\text{MMLM}} + \lambda \|\Theta\|^2, \quad (5)$$

where Θ contains all trainable parameters of the model, and λ is the coefficient of L_2 regularizer.

Tasks	Datasets	Train	Valid	Test
Natural Language Code Retrieval	AdvTest	251K	9.6K	19.2K
	CosQA	19.6K	0.5K	0.5K
	CoNaLa	2.4K	-	0.5K
Code-to-Code Retrieval	Python800	72K	4K	4K
Code Clone Detection	Python800	144K	8K	8K
Code Defect Detection	GREAT	100K	5K	5K

Table 2: Statistics of datasets for downstream tasks.

4 Experimental Setup

We conduct experiments to answer the following research questions: (1) *How effective is CODE-MVP compared with the state-of-the-art baselines?* (2) *How do different components and different views affect our CODE-MVP?*

4.1 Pre-Training Dataset and Settings

Different programming languages often require different program analyzers. Existing program analysis tools rarely support multiple programming languages and multi-view program transformations. For convenience, we choose Python for our experiments, as it is very popular and used in many projects. We pre-train CODE-MVP on the Python corpus of CodeSearchNet dataset (Husain et al., 2019), which consists of 0.5M bimodal Python functions with their corresponding natural-language comments, as well as 1.1M unimodal Python functions.

CODE-MVP is built on the top of Transformer (Vaswani et al., 2017), and consists of a 12-layer encoder with 768 hidden sizes and 12 attention heads. The pre-training procedure is conducted on 8 NVIDIA V100 GPUs for 600K steps, with each mini-batch containing 128 sequences up to 512 tokens including special tokens. According to the length distribution of samples in the training corpus, we set the lengths of PL/AST/CFG/PT in unpaired data to 512, and set the lengths of NL and PL/AST/CFG/PT in paired data to 96 and 416 respectively. The learning rate of CODE-MVP is set to $1e-4$ with a linear warm up over the first 30K steps and a linear decay. CODE-MVP is trained with a dropout rate of 0.1 on all layers and attention weights. We initialize the parameters of CODE-MVP by GraphCodeBERT (Guo et al., 2021) and utilize a BPE tokenizer (Sennrich et al., 2016).

4.2 Evaluation Tasks, Datasets and Metrics

We select several program comprehension tasks to evaluate CODE-MVP, including natural language code retrieval, code similarity, and code defect detection. We pre-train CODE-MVP on Python corpus, and choose several public Python datasets to evaluate it, as shown in Table 2.

Natural Language Code Retrieval. This task aims to find the most relevant code snippet from a collection of candidates, given a natural language query. We choose three datasets to evaluate this task, including AdvTest (Lu et al., 2021), CoNaLa (Yin et al., 2018), and CoSQA (Huang et al., 2021). We adopt the Mean Reciprocal Rank (MRR) metric to evaluate the performance of code retrieval. In AdvTest dataset, we set the learning rate as $5e-5$, the batch size as 32, the maximum fine-tuning epoch as 20, the maximum length of both query and code sequence as 256. In CoNaLa and CoSQA datasets, we set the learning rate as $5e-5$, the batch size as 32, the maximum fine-tuning epoch as 30, the maximum length of query and code sequence as 128. In AdvTest and CoSQA datasets, we save the optimal checkpoint on the validation set, and test it on the testing set. In CoNaLa dataset, we report the best results on the testing set.

Code Similarity. This task is always categorized into two groups: code-to-code retrieval and code clone detection. We conduct experiments on the Python800 dataset (Puri et al., 2021), which is composed of 800 problems with each problem having 300 unique Python solution files. We remove those files not in UTF-8 encoding formats and randomly select 100 solutions for each problem. In code-to-code retrieval, the filtered dataset is split to 720/40/40 problems for training, validation, and testing. Given a program, this task aims to retrieve other programs that solve the same problem; we evaluate using Mean Average Precision (MAP). Regarding the task of code clone detection, we treat it as binary classification and evaluate it using the Accuracy score, following (Puri et al., 2021).

To train these two tasks, we set the learning rate as $2e-5$, the batch size as 32, the epoch number as 20. In code-to-code retrieval, we set the maximum length of both query and code sequence as 256. In code clone detection, we set the maximum concatenation sequence length of the two code snippets to 512. We save the optimal checkpoint on the validation set, and test it on the testing set.

Models	AdvTest	CoNaLa	CoSQA	Average
RoBERTa	18.3	30.7	57.6	35.5
CodeBERT	27.2	38.9	64.2	43.4
GraphCodeBERT	35.2	47.3	68.2	50.2
PLBART	34.3	45.5	65.3	48.4
CodeT5	36.5	47.7	67.7	50.6
SynCoBERT	38.1	48.4	69.6	52.0
CODE-MVP	40.4	50.6	72.1	54.4

Table 3: Results on the natural language code retrieval task evaluating with MRR, using the AdvTest, CoNaLa, and CoSQA datasets.

Code Defect Detection. This task aims to identify whether a given piece of code snippet is vulnerable or not, which is usually treated as a binary classification task. We evaluate all models on the GREAT dataset (Hellendoorn et al., 2020), which is originally built from the ETH Py150 dataset (Raychev et al., 2016). We evaluate the performance of code defect detection using the Accuracy score. We randomly select 100K samples for training, 5K samples for validation and 5K samples for testing, respectively. We set the learning rate as $5e-5$, the batch size as 32, the maximum fine-tuning epoch as 50, the maximum length of both query and code sequence as 256. We save the optimal checkpoint on the validation set, and test it on the testing set.

4.3 Baselines

We compare CODE-MVP with various state-of-the-art models. **RoBERTa** (Liu et al., 2019) is a robustly optimized BERT (Devlin et al., 2019), which is originally pre-trained on a large-scale natural-language corpus. We fine-tune it on source code datasets of downstream tasks. **CodeBERT** (Feng et al., 2020) is pre-trained on NL-PL pairs using both masked language modeling (Devlin et al., 2019) and replaced token detection (Clark et al., 2020) objectives. **GraphCodeBERT** (Guo et al., 2021) is a pre-trained language model of source code which incorporates the data flow information of source code. **PLBART** (Ahmad et al., 2021) is based on the BART (Lewis et al., 2020) architecture and pre-trained on Python and Java functions using denoising autoencoding. **CodeT5** (Wang et al., 2021b) is based on the T5 (Raffel et al., 2020) architecture and employs denoising sequence-to-sequence pre-training on seven programming languages. **SynCoBERT** (Wang et al., 2021) incorporates AST by edge prediction and uses contrastive learning to maximize the mutual information among programs, documents, and ASTs.

Models	MAP@R	Accuracy
RoBERTa	82.9	94.4
CodeBERT	86.1	95.2
GraphCodeBERT	88.8	95.9
PLBART	86.7	95.5
CodeT5	88.1	95.7
SynCoBERT	89.2	96.1
CODE-MVP	91.5	97.4

Table 4: Results on the code-to-code retrieval and code clone detection tasks evaluating with MAP and Accuracy score, using the Python800 dataset.

5 Results and Analysis

5.1 Performance on Downstream Tasks (RQ1)

Natural Language Code Retrieval. Table 3 shows the results of natural language code retrieval on three datasets. We can observe that CODE-MVP outperforms all baseline models on all datasets. Specifically, it outperforms CodeT5 by 3.8 points on average. Compared to the previous state-of-the-art SynCoBERT, CODE-MVP also performs better with an average improvement of 2.4 points. This significant performance improvement indicates that the code representation learned by CODE-MVP preserves more code semantics. We attribute this improvement to our introduced multi-view contrastive pre-training strategy.

Code Similarity. Table 4 presents the results for code similarity calculation, including code-to-code retrieval and code clone detection. We can see that CODE-MVP significantly outperforms all baseline models on these two tasks. In the task of code-to-code retrieval, CODE-MVP outperforms CodeT5 and SynCoBERT by 3.4 points and 2.3 points, respectively. In the task of code clone detection, CODE-MVP achieves 1.5 and 1.3 points higher compared to GraphCodeBERT and SynCoBERT, respectively. These results show that CODE-MVP can better identify those programs with the same semantics and distinguish those programs with different semantics.

Code Defect Detection. Table 5 shows the experimental results of code defect detection. CODE-MVP consistently outperforms all models. Specifically, it outperforms GraphCodeBERT and SynCoBERT by 1.8 and 1.1 points, respectively. These results indicate that CODE-MVP can effectively preserve the semantics of programs, which is beneficial for code defect detection.

Models	Accuracy
RoBERTa	81.9
CodeBERT	85.5
GraphCodeBERT	87.5
PLBART	86.8
CodeT5	87.4
SynCoBERT	88.2
CODE-MVP	89.3

Table 5: Results on the code defect detection task evaluating with Accuracy score, using the GREAT dataset.

Models	AdvTest	CoNaLa	CoSQA	Average
CODE-MVP	40.4	50.6	72.1	54.4
w/o MVCL	36.2	47.7	69.2	51.0
w/o FGTI	38.0	48.9	70.8	52.6
w/o AST	39.1	48.5	71.3	53.0
w/o PT	38.2	48.6	70.8	52.5
w/o CFG	37.8	47.9	70.5	52.1

Table 6: Ablation study on the task of natural language code retrieval, evaluated using MRR.

5.2 Ablation Study (RQ2)

We empirically study several simplified variants of CODE-MVP to understand the contributions of each component, including the Multi-View Contrastive Learning (MVCL), Fine-Grained Type Inference (FGTI), Abstract Syntax Tree (AST), Program Transformation (PT), and Control Flow Graph (CFG). Taking the natural language code retrieval task as an example, Table 6 shows the experimental results of each variant on that task. The setting of w/o (MVCL, FGTI) indicates that these pre-training objectives are removed from CODE-MVP respectively. The setting of w/o (AST, PT, CFG) indicates that different views of programs are removed from CODE-MVP respectively. From Table 6, several meaningful observations can be drawn. (1) Both MVCL and FGTI effectively increase the performance, which confirms that the two proposed pre-training objectives can indeed improve the ability of the model for program comprehension. (2) Exploiting different views of programs can bring performance improvements to the model as arbitrarily discarding any view of programs degrades the performance. Additionally, the introduction of CFG brings more performance improvements, indicating the importance of execution information for program understanding.

6 Related Work

Pre-Trained Models for Source Code. Benefiting from the strong power of pre-trained models

in natural language processing (Liu et al., 2019; Devlin et al., 2019; Wang et al., 2021a, 2020a,b), several recent works attempt to use the pre-training techniques on programs (Svyatkovskiy et al., 2020). Kanade et al. (2020) proposed CuBERT which follows the architecture of BERT (Devlin et al., 2019), and is pre-trained with a masked language modeling objective on a large-scale Python corpus. Feng et al. (2020) proposed CodeBERT, which is pre-trained on NL-PL pairs in six programming languages, introducing the replaced token detection objective (Clark et al., 2020). Furthermore, Guo et al. (2021) proposed GraphCodeBERT, which incorporates the data flow of programs into the model pre-training process. Wang et al. (2021) proposed SynCoBERT, which incorporates ASTs via edge prediction to enhance the structural information of programs. They also used contrastive learning to maximize the mutual information among programs, documents, and ASTs. Lu et al. (2021) proposed CodeGPT for code completion, which is pre-trained using a unidirectional language modeling objective. Ahmad et al. (2021) proposed PLBART based on BART (Lewis et al., 2020), which is pre-trained on a large-scale corpus of Java and Python programs paired with their corresponding comments via denoising autoencoding. Wang et al. (2021b) proposed CodeT5 following the architecture of T5 (Raffel et al., 2020). It employs denoising sequence-to-sequence pre-training on seven programming languages. Recently, Wan et al. (2022b) conducted a thorough structural analysis aiming to provide an interpretation of pre-trained language models for source code (e.g., CodeBERT and GraphCodeBERT).

Program Analysis for Code Intelligence. In addition to the lexical information of programs, many recent works attempt to leverage program analysis techniques to capture the structural and syntactic representations of programs (Cummins et al., 2020). Kim et al. (2021) designed several strategies to feed the ASTs of programs into Transformer (Vaswani et al., 2017). Li et al. (2019) proposed a graph matching network, which utilizes the CFG of the program to deal with the challenge of binary function similarity search. Ling et al. (2021) proposed a deep graph matching and searching model based on graph neural networks (Kipf and Welling, 2017; Wang et al., 2021b,a; Yu et al., 2022; Zhao et al., 2022) for code retrieval. They represented both natural language queries and code snippets based on

the unified graph-structured data. Iyer et al. (2020) presented the program-derived semantic graph to capture the semantics of programs at multiple levels of abstraction. Ben-Nun et al. (2018) presented `inst2vec`, which locally embeds individual statement in LLVM intermediate representations by processing a contextual flow graph with a context prediction objective (Mikolov et al., 2013).

Contrastive Learning on Programs. Recently, several attempts have been made to leverage contrastive learning for better code semantics. `ContraCode` (Jain et al., 2021) and `Corder` (Bui et al., 2021b) first utilized semantic-preserving program transformations such as identifier renaming, dead code insertion, to build positive instances. Then a contrastive learning objective is designed to maximize the mutual information among the positive and negative instances. Ding et al. (2021) presented a self-supervised pre-training technique called `BOOST` based on contrastive learning. They inject real-world bugs to build hard negative pairs. In `CODE-MVP`, we construct the positive pairs throughout the compilation process of programs, including lexical analysis, syntax analysis, semantic analysis, and static analysis. It is the first pre-trained model that integrates multi-views of programs for program comprehension.

7 Conclusion

In this paper, we have proposed `CODE-MVP`, a novel approach to represent the source code with multi-view contrastive pre-training learning. We extract multiple code views with compiler tools and learn the complement among them under a contrastive learning framework. We also propose a fine-grained type inference task in the pre-training process. Comprehensive experiments on three downstream tasks over five datasets verify the effectiveness of `CODE-MVP` when compared with several state-of-the-art baselines.

Acknowledgements

We would like to thank Gerasimos Lampouras and Ignacio Iacobacci from Huawei London Research Institute for their constructive comments on this paper. Jin Liu is supported by National Natural Science Foundation of China under Grant No. 61972290. Yao Wan is partially supported by National Natural Science Foundation of China under Grant No. 62102157. Hao Wu is supported by National Natural Science Foundation of

China under Grant No. 61962061, and partially supported by Yunnan Provincial Foundation for Leaders of Disciplines in Science and Technology (202005AC160005).

References

- Wasi Uddin 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, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 2655–2668. Association for Computational Linguistics.
- Miltiadis Allamanis, Earl T. Barr, Premkumar T. Devanbu, and Charles Sutton. 2018. [A survey of machine learning for big code and naturalness](#). *ACM Comput. Surv.*, 51(4):81:1–81:37.
- Uri Alon, Meital Zilberstein, Omer Levy, and Eran Yahav. 2019. [code2vec: learning distributed representations of code](#). *Proc. ACM Program. Lang.*, 3(POPL):40:1–40:29.
- Jong-hoon (David) An, Avik Chaudhuri, Jeffrey S. Foster, and Michael Hicks. 2011. [Dynamic inference of static types for ruby](#). In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, pages 459–472. ACM.
- Tal Ben-Nun, Alice Shoshana Jakobovits, and Torsten Hoefler. 2018. [Neural code comprehension: A learnable representation of code semantics](#). In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 3589–3601.
- Nghi D. Q. Bui, Yijun Yu, and Lingxiao Jiang. 2021a. [Infercode: Self-supervised learning of code representations by predicting subtrees](#). In *43rd IEEE/ACM International Conference on Software Engineering, ICSE 2021, Madrid, Spain, 22-30 May 2021*, pages 1186–1197. IEEE.
- Nghi D. Q. Bui, Yijun Yu, and Lingxiao Jiang. 2021b. [Self-supervised contrastive learning for code retrieval and summarization via semantic-preserving transformations](#). In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 511–521. ACM.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. 2020. [A simple framework for contrastive learning of visual representations](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR.

- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: pre-training text encoders as discriminators rather than generators](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Chris Cummins, Zacharias V. Fisches, Tal Ben-Nun, Torsten Hoefler, and Hugh Leather. 2020. [Programl: Graph-based deep learning for program optimization and analysis](#). *CoRR*, abs/2003.10536.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Yanguibo Ding, Luca Buratti, Saurabh Pujar, Alessandro Morari, Baishakhi Ray, and Saikat Chakraborty. 2021. [Contrastive learning for source code with structural and functional properties](#). *CoRR*, abs/2110.03868.
- 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](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 1536–1547. Association for Computational Linguistics.
- Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, Shujie Liu, Long Zhou, Nan Duan, Alexey Svyatkovskiy, Shengyu Fu, Michele Tufano, Shao Kun Deng, Colin B. Clement, Dawn Drain, Neel Sundaresan, Jian Yin, Daxin Jiang, and Ming Zhou. 2021. [Graphcodebert: Pre-training code representations with data flow](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Vincent J. Hellendoorn, Charles Sutton, Rishabh Singh, Petros Maniatis, and David Bieber. 2020. [Global relational models of source code](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Junjie Huang, Duyu Tang, Linjun Shou, Ming Gong, Ke Xu, Daxin Jiang, Ming Zhou, and Nan Duan. 2021. [Cosqa: 20, 000+ web queries for code search and question answering](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 5690–5700. Association for Computational Linguistics.
- Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. 2019. [Code-searchnet challenge: Evaluating the state of semantic code search](#). *CoRR*, abs/1909.09436.
- Roshni G. Iyer, Yizhou Sun, Wei Wang, and Justin Gottschlich. 2020. [Software language comprehension using a program-derived semantic graph](#). *CoRR*, abs/2004.00768.
- 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, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Paras Jain, Ajay Jain, Tianjun Zhang, Pieter Abbeel, Joseph Gonzalez, and Ion Stoica. 2021. [Contrastive code representation learning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 5954–5971. Association for Computational Linguistics.
- Faizan Javed, Barrett R. Bryant, Matej Crepinsek, Marjan Mernik, and Alan P. Sprague. 2004. [Context-free grammar induction using genetic programming](#). In *Proceedings of the 42nd Annual Southeast Regional Conference, 2004, Huntsville, Alabama, USA, April 2-3, 2004*, pages 404–405. ACM.
- Xue Jiang, Zhuoran Zheng, Chen Lyu, Liang Li, and Lei Lyu. 2021. [Treebert: A tree-based pre-trained model for programming language](#). In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI 2021, Virtual Event, 27-30 July 2021*, volume 161 of *Proceedings of Machine Learning Research*, pages 54–63. AUAI Press.
- Aditya Kanade, Petros Maniatis, Gogul Balakrishnan, and Kensen Shi. 2020. [Learning and evaluating contextual embedding of source code](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 5110–5121. PMLR.
- Seohyun Kim, Jinman Zhao, Yuchi Tian, and Satish Chandra. 2021. [Code prediction by feeding trees to transformers](#). In *43rd IEEE/ACM International Conference on Software Engineering, ICSE 2021, Madrid, Spain, 22-30 May 2021*, pages 150–162. IEEE.
- Thomas N. Kipf and Max Welling. 2017. [Semi-supervised classification with graph convolutional networks](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.

- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7871–7880. Association for Computational Linguistics.
- Li Li, Jiawei Wang, and Haowei Quan. 2022. [Scalpel: The python static analysis framework](#). *CoRR*, abs/2202.11840.
- Yujia Li, Chenjie Gu, Thomas Dullien, Oriol Vinyals, and Pushmeet Kohli. 2019. [Graph matching networks for learning the similarity of graph structured objects](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 3835–3845. PMLR.
- Xiang Ling, Lingfei Wu, Saizhuo Wang, Gaoning Pan, Tengfei Ma, Fangli Xu, Alex X. Liu, Chunming Wu, and Shouling Ji. 2021. [Deep graph matching and searching for semantic code retrieval](#). *ACM Trans. Knowl. Discov. Data*, 15(5):88:1–88:21.
- 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](#). *CoRR*, abs/1907.11692.
- Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin B. Clement, Dawn Drain, Daxin Jiang, Duyu Tang, Ge Li, Lidong Zhou, Linjun Shou, Long Zhou, Michele Tufano, Ming Gong, Ming Zhou, Nan Duan, Neel Sundaresan, Shao Kun Deng, Shengyu Fu, and Shujie Liu. 2021. [Codexglue: A machine learning benchmark dataset for code understanding and generation](#). *CoRR*, abs/2102.04664.
- Charith Mendis, Alex Renda, Saman P. Amarasinghe, and Michael Carbin. 2019. [Ithemal: Accurate, portable and fast basic block throughput estimation using deep neural networks](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 4505–4515. PMLR.
- Tomás Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 3111–3119.
- Safa Omri and Carsten Sinz. 2020. [Deep learning for software defect prediction: A survey](#). In *ICSE '20: 42nd International Conference on Software Engineering, Workshops, Seoul, Republic of Korea, 27 June - 19 July, 2020*, pages 209–214. ACM.
- Jukka Paakki. 1995. [Attribute grammar paradigms - A high-level methodology in language implementation](#). *ACM Comput. Surv.*, 27(2):196–255.
- Long N. Phan, Hieu Tran, Daniel Le, Hieu Nguyen, James T. Anibal, Alec Peltekian, and Yanfang Ye. 2021. [Cotext: Multi-task learning with code-text transformer](#). *CoRR*, abs/2105.08645.
- Ruchir Puri, David S. Kung, Geert Janssen, Wei Zhang, Giacomo Domeniconi, Vladimir Zolotov, Julian Dolby, Jie Chen, Mihir R. Choudhury, Lindsey Decker, Veronika Thost, Luca Buratti, Saurabh Pujar, and Ulrich Finkler. 2021. [Project codenet: A large-scale AI for code dataset for learning a diversity of coding tasks](#). *CoRR*, abs/2105.12655.
- Md. Rafiqul Islam Rabin, Nghi D. Q. Bui, Yijun Yu, Lingxiao Jiang, and Mohammad Amin Alipour. 2020. [On the generalizability of neural program analyzers with respect to semantic-preserving program transformations](#). *CoRR*, abs/2008.01566.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Veselin Raychev, Pavol Bielek, and Martin T. Vechev. 2016. [Probabilistic model for code with decision trees](#). In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2016, part of SPLASH 2016, Amsterdam, The Netherlands, October 30 - November 4, 2016*, pages 731–747. ACM.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). *ArXiv*, abs/1508.07909.
- Alexey Svyatkovskiy, Shao Kun Deng, Shengyu Fu, and Neel Sundaresan. 2020. [Intellicode compose: code generation using transformer](#). In *ESEC/FSE '20: 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Virtual Event, USA, November 8-13, 2020*, pages 1433–1443. ACM.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Yao Wan, Yang He, Zhangqian Bi, Jianguo Zhang, Yulei Sui, Hongyu Zhang, Kazuma Hashimoto, Hai Jin, Guandong Xu, Caiming Xiong, and Philip S. Yu. 2022a. [Naturalcc: An open-source toolkit for code](#)

- intelligence. In *Proceedings of 44th International Conference on Software Engineering, Companion Volume*. ACM.
- Yao Wan, Jingdong Shu, Yulei Sui, Guandong Xu, Zhou Zhao, Jian Wu, and Philip S. Yu. 2019. **Multi-modal attention network learning for semantic source code retrieval**. In *34th IEEE/ACM International Conference on Automated Software Engineering, ASE 2019, San Diego, CA, USA, November 11-15, 2019*, pages 13–25. IEEE.
- Yao Wan, Wei Zhao, Hongyu Zhang, Yulei Sui, Guandong Xu, and Hai Jin. 2022b. **What do they capture?—a structural analysis of pre-trained language models for source code**. In *Proceedings of the 44th International Conference on Software Engineering*.
- Yao Wan, Zhou Zhao, Min Yang, Guandong Xu, Haochao Ying, Jian Wu, and Philip S Yu. 2018. **Improving automatic source code summarization via deep reinforcement learning**. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, pages 397–407.
- Ke Wang and Zhendong Su. 2020. **Blended, precise semantic program embeddings**. In *Proceedings of the 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2020, London, UK, June 15-20, 2020*, pages 121–134. ACM.
- Xin Wang, Jin Liu, Li Li, Xiao Chen, Xiao Liu, and Hao Wu. 2020a. **Detecting and explaining self-admitted technical debts with attention-based neural networks**. In *35th IEEE/ACM International Conference on Automated Software Engineering, ASE 2020, Melbourne, Australia, September 21-25, 2020*, pages 871–882. IEEE.
- Xin Wang, Jin Liu, Xiao Liu, Xiaohui Cui, and Hao Wu. 2020b. **A spatial and sequential combined method for web service classification**. In *Web and Big Data - 4th International Joint Conference, APWeb-WAIM 2020, Tianjin, China, September 18-20, 2020, Proceedings, Part I*, volume 12317 of *Lecture Notes in Computer Science*, pages 764–778. Springer.
- Xin Wang, Xiao Liu, Li Li, Xiao Chen, Jin Liu, and Hao Wu. 2021a. **Time-aware user modeling with check-in time prediction for next POI recommendation**. In *2021 IEEE International Conference on Web Services, ICWS 2021, Chicago, IL, USA, September 5-10, 2021*, pages 125–134. IEEE.
- Xin Wang, Xiao Liu, Jin Liu, and Hao Wu. 2021b. **Relational graph neural network with neighbor interactions for bundle recommendation service**. In *2021 IEEE International Conference on Web Services, ICWS 2021, Chicago, IL, USA, September 5-10, 2021*, pages 167–172. IEEE.
- Xin Wang, Yasheng Wang, Fei Mi, Pingyi Zhou, Yao Wan, Xiao Liu, Li Li, Hao Wu, Jin Liu, and Xin Jiang. 2021. **Syncobert: Syntax-guided multi-modal contrastive pre-training for code representation**. Xin Wang, Yasheng Wang, Yao Wan, Fei Mi, Yitong Li, Pingyi Zhou, Jin Liu, Hao Wu, Xin Jiang, and Qun Liu. 2022. **Compilable neural code generation with compiler feedback**. volume abs/2203.05132.
- Xin Wang, Pingyi Zhou, Yasheng Wang, Xiao Liu, Jin Liu, and Hao Wu. 2021a. **Servicebert: A pre-trained model for web service tagging and recommendation**. In *Service-Oriented Computing - 19th International Conference, ICSOC 2021, Virtual Event, November 22-25, 2021, Proceedings*, volume 13121 of *Lecture Notes in Computer Science*, pages 464–478. Springer.
- Yue Wang, Weishi Wang, Shafiq R. Joty, and Steven C. H. Hoi. 2021b. **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, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 8696–8708. Association for Computational Linguistics.
- Martin White, Michele Tufano, Christopher Vendome, and Denys Poshyvanyk. 2016. **Deep learning code fragments for code clone detection**. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering, ASE 2016, Singapore, September 3-7, 2016*, pages 87–98. ACM.
- Hao Wu, Yunhao Duan, Kun Yue, and Lei Zhang. 2021. **Mashup-oriented web api recommendation via multi-model fusion and multi-task learning**. *IEEE Transactions on Services Computing*, pages 1–1.
- Pengcheng Yin, Bowen Deng, Edgar Chen, Bogdan Vasilescu, and Graham Neubig. 2018. **Learning to mine aligned code and natural language pairs from stack overflow**. In *Proceedings of the 15th International Conference on Mining Software Repositories, MSR 2018, Gothenburg, Sweden, May 28-29, 2018*, pages 476–486. ACM.
- Jiaojiao Yu, Kunsong Zhao, Jin Liu, Xiao Liu, Zhou Xu, and Xin Wang. 2022. **Exploiting gated graph neural network for detecting and explaining self-admitted technical debts**. *J. Syst. Softw.*, 187:111219.
- Kunsong Zhao, Jin Liu, Zhou Xu, Xiao Liu, Lei Xue, Zhiwen Xie, Yuxuan Zhou, and Xin Wang. 2022. **Graph4web: A relation-aware graph attention network for web service classification**. *Journal of Systems and Software*, 190:111324.
- Kunsong Zhao, Zhou Xu, Meng Yan, Tao Zhang, Dan Yang, and Wei Li. 2021a. **A comprehensive investigation of the impact of feature selection techniques on crashing fault residence prediction models**. *Inf. Softw. Technol.*, 139:106652.
- Kunsong Zhao, Zhou Xu, Tao Zhang, Yutian Tang, and Meng Yan. 2021b. **Simplified deep forest model based just-in-time defect prediction for android mobile apps**. *IEEE Trans. Reliab.*, 70(2):848–859.