

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

5-2022

On the transferability of pre-trained language models for low-resource programming languages

Fuxiang CHEN

Fatemeh H. FARD

David LO

Singapore Management University, davidlo@smu.edu.sg

Timofey BRYKSIN

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#)

Citation

CHEN, Fuxiang; FARD, Fatemeh H.; LO, David; and BRYKSIN, Timofey. On the transferability of pre-trained language models for low-resource programming languages. (2022). *Proceedings of the 30th IEEE/ACM International Conference on Program Comprehension, ICPC 2022, Virtual Event, May 16-17*. 401-412. Available at: https://ink.library.smu.edu.sg/sis_research/7693

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.

On the Transferability of Pre-trained Language Models for Low-Resource Programming Languages

Fuxiang Chen
University of British Columbia
Canada
fuxiang.chen@ubc.ca

David Lo
Singapore Management University
Singapore
davidlo@smu.edu.sg

Fatemeh H. Fard
University of British Columbia
Canada
fatemeh.fard@ubc.ca

Timofey Bryksin
JetBrains Research
Republic of Cyprus
timofey.bryksin@jetbrains.com

ABSTRACT

A recent study by Ahmed and Devanbu reported that using a corpus of code written in multilingual datasets to *fine-tune* multilingual Pre-trained Language Models (PLMs) achieves higher performance as opposed to using a corpus of code written in just one programming language. However, no analysis was made with respect to fine-tuning monolingual PLMs. Furthermore, some programming languages are inherently different and code written in one language usually cannot be interchanged with the others, i.e., Ruby and Java code possess very different structure. To better understand how monolingual and multilingual PLMs affect different programming languages, we investigate 1) the performance of PLMs on Ruby for two popular Software Engineering tasks: Code Summarization and Code Search, 2) the strategy (to select programming languages) that works well on fine-tuning multilingual PLMs for Ruby, and 3) the performance of the fine-tuned PLMs on Ruby given different code lengths.

In this work, we analyze over a hundred of pre-trained and fine-tuned models. Our results show that 1) multilingual PLMs have a lower Performance-to-Time Ratio (the BLEU, METEOR, or MRR scores over the fine-tuning duration) as compared to monolingual PLMs, 2) our proposed strategy to select target programming languages to fine-tune multilingual PLMs is effective – it reduces the time to fine-tune yet achieves higher performance in Code Summarization and Code Search tasks, and 3) our proposed strategy consistently shows good performance on different code lengths.

CCS CONCEPTS

• **Software and its engineering** → **General programming languages**; • **Computing methodologies** → *Artificial intelligence*.

KEYWORDS

pre-trained language models, low-resource languages

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 ACM 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.

ICPC '22, May 16–17, 2022, Virtual Event, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

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

ACM Reference Format:

Fuxiang Chen, Fatemeh H. Fard, David Lo, and Timofey Bryksin. 2022. On the Transferability of Pre-trained Language Models for Low-Resource Programming Languages. In *30th International Conference on Program Comprehension (ICPC '22)*, May 16–17, 2022, Virtual Event, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Unsupervised pre-training of language models on large corpora significantly improves the performance in many downstream tasks [11, 30]. In this work, we refer to a Pre-trained Language Model and its plural form as PLM and PLMs, respectively. There have been some attempts to understand how PLMs affect the performance of different downstream tasks empirically [4, 8, 12, 13, 31, 33].

Despite the existing efforts to understand PLMs, there are still many unknowns on the transferability of PLMs for programming languages. **Firstly**, existing PLMs are trained either on a programming language or on multiple programming languages. Little is known if PLMs trained on a particular programming language yields better performance than a more general PLM that is trained on multiple programming languages. A closely related work is by Ahmed and Devanbu [3] that studied the effects of *fine-tuning* for publicly available multilingual pre-trained models, CodeBERT and GraphCodeBERT. However, some programming languages are inherently different, so utilizing a single multilingual model may not always yield the best performance. **Secondly**, the zero-shot setting in previous work is lightly studied [12]. For example, other programming languages may have very different structures as compared to the programming languages used to pre-train PLMs. Also, during the pre-training, fine-tuning, and testing processes, different programming languages may be used. To better understand the zero-shot setting, these need to be studied more thoroughly. **Thirdly**, the current datasets that are released for training PLMs on Software Engineering (SE) related tasks such as Code Summarization, are dominated by a few languages, mainly Java and Python [17] – they are known as high-resource programming languages as there is a high volume of code written in them. Other programming languages are often missing or have low number of records in the datasets – they are known as low-resource programming languages [17]. In a recent survey conducted by StackOverflow, we observe that although Java and Python are among the popular programming languages used by developers, developers also reported

that they are using 36 other programming languages such as Ruby, Kotlin, and Scala [1]. Thus, understanding if PLMs pre-trained on high-resource programming languages can be utilized for other programming languages is important.

To bridge the gap in understanding the applicability and the transferability of PLMs in SE, in this study, **we focus on studying the impact of PLMs on a low-resource programming language corpus — specifically, we choose Ruby as the study subject because it is highly ranked among low-resource languages in the Stack Overflow survey [1] and it is also a commonly used low-resource programming language [12, 15, 17].** We explore five different settings when using PLMs for different downstream tasks:

- (1) Transferability of PLMs pre-trained with a code base written in a single and multiple programming languages: we are interested to know if a monolingual PLM (*a PLM that is pre-trained and fine-tuned on a single programming language*) works better than a multilingual PLM (CodeBERT and GraphCodeBERT).
- (2) Transferability of PLMs in different zero-shot scenarios: we are interested to know if unseen programming languages can leverage a PLM effectively in downstream tasks.
- (3) Efficiency of PLMs: we are interested to know the performance and training time trade-offs among the PLMs.
- (4) Transferability of PLMs depending on different code lengths: as developers write code differently, to better understand how a PLM may perform on code of different lengths, in the fourth setting, we are interested to know the performance of PLMs on different code length.
- (5) The strategy to select suitable programming languages for fine-tuning: we are interested to explore a strategy to select suitable high-resource programming languages for fine-tuning.

For the dataset, we use CodeSearchNet [17] which contains code in six different programming languages, including high-resource and low-resource ones. We note here that the purpose of this study is not to beat the state-of-the-art, but to understand the things mentioned above. Thus, RoBERTa, a strong baseline model used in many PLM related studies and is the basis of the multilingual PLMs (CodeBERT and GraphCodeBERT), is used here for pre-training the PLMs [9, 12, 19, 22, 36, 37, 48]. For all the settings, over 100 PLMs are trained or fine-tuned on different programming languages, and we evaluate their performance on two commonly studied SE tasks: Code Summarization and Code Search [28, 43, 45, 49].

Our results show several interesting phenomena: 1) For Code Summarization, PLMs fine-tuned on the entire multilingual dataset do not yield the best performance but, for Code Search, the best performance is observed on PLMs fine-tuned on the entire multilingual dataset; 2) Monolingual PLMs trained on a combined multilingual dataset have higher Performance-to-Time Ratio (PTR) than multilingual PLMs trained on a multilingual dataset. The PTR ratio measures the trade-off between the training time to fine-tune a PLM and its performance in downstream tasks; 3) PLMs fine-tuned on the Python dataset have the best performance in our zero-shot experiments; 4) There are negligible differences in the performance between PLMs tested on a test dataset binned in different code

lengths and PLMs tested on the entire test dataset; and 5) Our proposed strategy in selecting a programming language for fine-tuning is effective: it improves the performance over PLMs fine-tuned on the combined multilingual dataset. The findings in this study are important since researchers and practitioners can save time (through our PLM study) and may achieve better performance by making a more informed decision in using PLM for their tasks. We note that the most recent empirical work on PLMs [3] reported that multilingual PLMs fine-tuned on a combined multilingual dataset perform better in downstream tasks, but our experiments have shown that this might not be the case for all the downstream tasks and we proposed an effective strategy to pick another high-resource language to train on (rather than training on the combined multilingual dataset).

Overall, this paper makes the following contributions:

- **An empirical evaluation on the downstream tasks using monolingual and multilingual PLMs** We perform a detailed quantitative and qualitative evaluation for over a hundred of models on two downstream tasks (Code Summarization and Code Search).
- **Proposed Strategy to select a suitable PL for fine-tuning PLMs** We proposed an effective strategy to select a suitable programming language for fine-tuning PLMs in Code Summarization and Code Search.
- **Multiple PLMs were trained on different programming languages for two different tasks – Code Summarization and Code Search.** Training a PLM requires high computational resources. In order to understand the applicability and transferability of PLMs in SE, we have pre-trained and fine-tuned over a hundred PLMs¹ in different programming languages. Based on our findings, in the downstream tasks, developer can use monolingual PLMs on fine-tuning the combined multilingual datasets which is more time-efficient yet having similar (Code Summarization) or better (Code Search) performance.

The rest of this paper is organized as follows. Related work is surveyed and discussed in Section 2. Sections 3 and 4 describe our research questions and methods that we employ to answer them, whereas Section 5 describes the experimental setup. We portray the results for both Code Summarization and Code Search in Section 6. Further discussions on the results are presented in Section 7 and various threats are analysed in Section 8. We conclude with directions for future research in Section 9.

2 RELATED WORK

Here, we surveyed how different PLMs are used in Software Engineering and discussed the missing gaps.

Kanade et al. use BERT to pre-train a model on Python source code [21]. The authors later train a BERT model for source code, known as CuBERT, which is then fine-tuned for classification tasks (e.g., wrong binary operator) and program repair tasks [22]. SCELMO is a PLM based on ELMO that is trained on JavaScript source code for the program repair task [23]. Xu et al. [47] incorporate external knowledge for code generation through pre-training the model with natural language and code pairs and then fine-tune

¹<https://doi.org/10.20383/102.0563>

their model for code generation. Buratti et al. train BERT on the C language, known as C-BERT, which is used for abstract syntax tree tagging tasks [7]. PyMT5 uses Transformer to pre-train a model for generating Python methods from docstrings. It also generates code summaries [10]. Pre-training a model to represent source code using contrastive learning is proposed in [7]. The authors present Corder and use this pre-trained model for Code Retrieval and Code Summarization tasks. GraphCodeBERT is a PLM that uses Transformer as its main architecture. The model is tested on Code Search, Code Clone Detection, Code Translation, and Code Refinement tasks [15]. CodeBERT is a PLM that uses programming and natural language in the pre-training and combines two training objectives; it is tested on different tasks such as Code Search [12]. IntelliCode Compose is a GPT based model that is trained on Python, C#, JavaScript, and TypeScript for code completion [41]. PLBART is another attempt to build a PLM using the BART architecture [24]. It is trained on Java and Python and is tested on several tasks including Code Summarization, Code Generation, Code Translation, Code Clone Detection, and Program Repair [2]. Text-To-Text Transfer Transformer is another study based on T5 [39] that leverages PLMs to study code related tasks including bug fixing and comment generation [32]. Other similar models that leverage PLMs for code related tasks are CodeT5 [46], Transcoder [40], SynCoBERT [44], TreeBERT [18], and a model that uses multi-task learning for pre-training the language model for code completion [27]. All of the above works present a PLM to represent code for different tasks. Some works focus on pre-training a PLM using a programming language, while others use two to six different programming languages. Little is known if PLMs that are pre-trained on a single programming language yield better results or generalize better on downstream tasks than those that are pre-trained on multiple programming languages.

Some recent works studied deep neural models leveraging PLMs empirically, mostly in their performance on downstream tasks. Ciniselli et al. [8] studied how RoBERTa and T5 models affect code completion. The authors concluded that Transformer-based models achieve good results when predicting an entire block of code or when predicting a few masked tokens. Mahmud et al. compare three Code Summarization models quantitatively and qualitatively [31]. Nie et al. [33] study the evaluation methodologies for comment generation and method namings, and proposed using time-segmented data for more realistic evaluation, i.e., take train, validation, and test sets from the same year range. The metrics, datasets used, and evaluation of comment generation models are also explored in the work of Gros et al. [13]. In another study, the limitations of large language models for program synthesis are explored [4]. A more related work by Feng et al. has shown the performance of Code Summarization on C# code leveraging PLMs in a zero-shot setting [12]. However, the zero-shot experiment is not the focus of their study and it was conducted very briefly. Little is known on the PLM's ability for zero-shot learning or on the transferability to different programming languages. Although these works explore different PLMs or models developed for a specific code related task, none of them probes the capability of a PLM in different settings, which are currently missing in the literature. With this study, we intend to start filling this gap. It is important to understand that so that researchers and practitioners can be more productive and

effective by making more informed decision in using PLMs for their tasks.

3 RESEARCH METHOD

To better understand the applicability and transferability of the PLMs in SE, we investigate a number of Research Questions (Section 3.1) and describe their study design (Section 3.2).

3.1 Research Questions

RQ1: Does training and fine-tuning on the individual programming languages improve the performance over multilingual PLMs that are fine-tuned on multilingual datasets?

Existing work have used either a single programming language or multiple programming languages (multilingual) to pre-train PLMs. Specially, the multilingual PLMs are CodeBERT and GraphCodeBERT. We are interested to know if additional programming languages add merits to pre-training a PLM in different tasks.

RQ2: Which PLM has the best Performance-to-Time Ratio (PTR)? Although the performance of a PLM is important, the training of a PLM is notoriously known to be computationally expensive. We are interested to know the PLMs that have the best trade-off between performance and training time.

RQ3: What are the best settings for zero-shot downstream tasks? Feng et al. has conducted a small study on zero-shot Code Summarization using CodeBERT [12]. The unseen programming language used in their experiment was C# and it has a similar structure to Java, which was used as one of the programming languages to pre-train CodeBERT. There are multiple unknowns in understanding the zero-shot settings on PLMs. For example, if we fine-tune a PLM (pre-trained with programming language A, e.g., PHP) on a programming language B (e.g., Go) and test it on a programming language C (e.g., Ruby), will it have a better performance over a PLM pre-trained and fine-tuned using the same target programming language? We are interested to understand this zero-shot setting.

RQ4: What effect does the PLMs have on different code length? Developers write code in different lengths. However, in many of the existing studies, the performance of the PLM is reported as an average metric score. Thus, it is possible that the reported score may be skewed towards certain code length within the test data — for example, the test data may contain mostly short code and the reported score may not reflect the true behavior of code in other lengths. We are interested to understand if the PLM has any effect on the length of the code.

RQ5: How effective is our strategy to decide in advance a language that can work well for a target low-resource language? Different programming languages have different syntax, and code fragments written in different languages are usually non-interchangeable. For example, code written in Ruby and Java have several differences such as the way data is flowed within code, and we cannot just rewrite a Ruby fragment with Java constructs and expect it to work. Here, we propose a strategy to choose in advance

programming languages that can work well for fine-tuning multilingual PLMs. We are interested to understand if our proposed strategy is effective.

3.2 Study Design

RQ1 Design: We train multiple PLMs that will be used for fine-tuning on two different widely used downstream tasks: Code Summarization and Code Search [3, 12, 15]. The PLMs are pre-trained on individual programming languages. We use CodeSearchNet, a popular dataset consisting of six programming languages published by Github and Microsoft [17] to pre-train and fine-tune PLMs. We compare *monolingual PLMs* (pre-trained on individual programming languages) fine-tuned on a *monolingual dataset* with the PLMs fine-tuned on a *multilingual datasets*. Here, the multilingual dataset refers to the combined dataset of all the programming languages. We also compare the best monolingual PLMs fine-tuned on each monolingual dataset with multilingual PLMs (CodeBERT and GraphCodeBERT) fine-tuned on the multilingual dataset which have reported having the best performance in Code Summarization and Code Search [3]. Additionally, we perform human evaluation on the Code Summarization task.

RQ2 Design: Here, we compute the Performance-to-Time Ratio – we measure the training time it takes to fine-tune a PLM and compare the time with its performance in the downstream tasks. For Code Summarization, the performance is measured in BLEU and METEOR, while for Code Search, the performance is measured in MRR. The time and performance are all normalized within the range of 0 and 1 prior to computing the ratio. We normalized the training time by having the largest training time as the denominator of all training times – this will have the effect of the largest training time having a normalized value of 1 and other training times scaled with respect to it. We normalized the performance in a similar manner, except that instead of the training time, we use the performance metric scores: BLEU, METEOR, and MRR.

RQ3 Design: We compare among the PLMs that were not pre-trained nor fine-tuned using the target language. We first pre-train the monolingual PLMs (except for Ruby). Then, we fine-tuned the monolingual PLMs using different monolingual datasets (except for Ruby). For each monolingual PLM, there are five fine-tuned models (based on the five monolingual datasets). For each monolingual PLM, we compare with all its fine-tuned models to find the best performing one. Then, among the best fine-tuned models in all the PLMs, we compare them to find the best model.

RQ4 Design: We compare the effects on the different code lengths on Code Summarization, similar to previous work [25]. Additionally, we also compare the effects on the different code lengths on Code Search, which has not been studied. We segregate the target test data into four different code length groups, based on the length distribution of code fragments in the target programming language. The four groups are: 1) code length between 0 and first quartile, 2) code length between first and second quartile, 3) code length between second and third quartile, and 4) code length on and above third quartile. After segregating the test data into these four groups,

we test all PLMs on each group. Specifically, for each PLM, we test all the fine-tuned models on each group. Among them, we compare to find the best fine-tuned model. We then compare the best fine-tuned models of each PLM to find the overall best model. Additionally, we also compare with the multilingual PLMs fine-tuned with the combined dataset which has reported having good performance [3].

RQ5 Design We propose a strategy (Section 4) to select a set of programming languages that works well in multilingual PLMs. We fine-tuned the PLMs with this selected set of programming languages.

For all the RQs, we use CodeBERT and GraphCodeBERT, the two state-of-the-art PLMs for code, as our baseline. They are pre-trained and fine-tuned on the multilingual datasets.

4 SELECTING A PROGRAMMING LANGUAGE FOR FINE-TUNING A PLM

4.1 Code Summarization

Overview We consider a suitable programming language that can be used to fine-tune a PLM for a target programming language to have both similar semantics and textual properties compared to the target programming language. We first train an embedding model using the whole multilingual dataset. Then, we compute semantic similarity between the individual monolingual datasets and the target programming language dataset. Afterward, we compute textual similarity between monolingual datasets and the target programming language dataset. Finally, we select the suitable programming language based on our proposed formula, which takes into account both the semantic and textual similarity scores.

Semantic Similarity. To detect similar semantics between programming languages, we train an embedding model (using the whole multilingual dataset) before computing the cosine similarity between the programming language and the target programming language. We made use of a recent paragraph embedding model that has reported having good performance in computing similarities between sentences by training the word embeddings jointly with bigram and trigram embeddings [16]. To train the embedding model, for every code function, we remove all line breaks and treat the code as a sequence of continuous word tokens – this sequence of continuous word tokens can be viewed as a sequence of sentence. The trained embeddings model is then used to retrieve the embeddings of every monolingual code and the target programming language code for computing cosine similarity. Finally, between each programming language and the target programming language, we compute the average similarity score. We further normalized this score to be in the range of 0 and 1, by dividing all the values with the largest average similarity score.

Textual Similarity To detect similar text between a programming language and the target programming language, we made use of CCFinder [20], a token-based code clone detector to compute the textual similarity between the programming language and the target programming language. CCFinder is able to detect code clones in a variety of format, including C, C++, C#, Cobol, Java, VB and plaintext. As our purpose is to detect “textual similarity” between cross programming languages code, we have used the plaintext

mode in CCFinder. Here, the textual similarity is the number of code clones detected. CCFinder transforms code fragments into suffix trees and uses them to detect exact and near-miss code clones. To have a more fine-grained detection, we set the minimum number of detected tokens to be 30, instead of the default 50. We perform code clone detection between every programming language and the target programming language. Finally, between each programming language and the target programming language, the number of code clones is computed. Similarly, we normalized this number to be in the range of 0 and 1, by dividing all the values with the largest number of code clones.

Suitability Function The suitability of a programming language is then formulated as:

$$\frac{Sim_{sem} + Sim_{text}}{2} \geq \theta \quad (1)$$

where Sim_{sem} refers to the normalized cosine similarity between the programming language and the target programming language, and Sim_{text} refers to the normalized number of code clones between the programming language and the target programming language. Following previous work, we set θ to be 0.5 [26].

4.2 Code Search

Based on our empirical experiments, we have observed that PLMs fine-tuned with the combined multilingual dataset perform best in Code Search (Table 5). Thus, for the Code Search task, we propose using the combined multilingual dataset to fine-tune the PLMs.

5 EXPERIMENTAL SETUP

In this section, we describe in detail 1) the PLM and the dataset used for training the PLM, 2) the downstream tasks and the datasets used for fine-tuning the models for these downstream tasks, 3) the evaluation metrics for the downstream tasks, and 4) the qualitative analysis for Code Summarization in Section 5.1, 5.2, 5.3, and 5.4 respectively.

5.1 The PLM and the Pretraining Dataset

RoBERTa enhances the pre-training task of BERT [11] and achieved higher performance in many NLP tasks [30]. It uses a bidirectional Masked Language Modeling (MLM) objective, where the model is trained to predict the masked tokens in the input text. In MLM, a small number of words are masked (15%) and the model is trained to predict them [9]. We use RoBERTa as it is a common base model used in many PLM studies in SE such as CodeBERT and GraphCodeBERT [12, 15]. We are unable to use CodeBERT or GraphCodeBERT for pre-training from scratch as their source code is close-sourced and not available.

Table 1: Dataset used for training PLM.

Language	bimodal DATA	unimodal CODES
Go	317,832	726,768
Java	500,754	1,569,889
Javascript	143,252	1,857,835
PHP	662,907	977,821
Python	458,219	1,156,085
Ruby	52,905	164,048

Dataset. As shown in Table 1, we train RoBERTa using the CodeSearchNet data [17], a dataset published by GitHub and Microsoft. It contains two different types of data: 1) parallel data of natural language-code pairs, known as bimodal data (column two) and 2) codes without paired natural language and natural language without paired codes, known as unimodal data (column three). Each unimodal code is a function without paired documentation. The programming languages in this dataset are Go, Java, Javascript, PHP, Python and Ruby. There are 2.1M bimodal data points and 6.4M unimodal codes. This dataset is commonly used in previous studies [12, 15, 17].

5.2 Downstream Tasks and Datasets

Task #1: Code Summarization. The Code Summarization task is to generate textual summaries describing the code, where the input to the model is a code snippet and the output is a description of the code functionality in natural language. For fine-tuning, we followed the CodeBERT paper and used their published code – for Code Summarization, an encoder-decoder framework is used to train a model to generate summaries, while for Code Search, the representation of [CLS] is used to measure the semantic relevance between the code and query [12].

Dataset for Task #1. We leverage the same dataset as described in Section 5.1 but use only the code-comment pairs to fine-tune the Code Summarization models. The dataset is pre-processed by the publishers and the cleaning scripts are provided. We use the same cleaned dataset in our experiments. It is split into train, test, and validation sets, and we use the same split to train our models. Table 2 shows the statistics of the dataset. We note that Ruby is a low-resource programming language.

Table 2: Dataset for Code Summarization.

Language	Train	Valid	Test
Go	317,832	14,242	14,291
Java	454,451	15,328	26,909
Javascript	123,889	8,253	6,483
PHP	523,712	26,015	28,391
Python	412,178	23,107	22,176
Ruby	48,791	2,209	2,279

Task #2: Code Search. The Code Search task is to find the most semantically related code from a collection of codes, given a natural language as the input. For fine-tuning, we followed the CodeBERT paper [12].

Dataset for Task #2. We used a preprocessed version of the CodeSearchNet data where a correct pair of test data <docstring, code> is combined with a fixed set of 999 incorrect pairs of test data [12]. This data preprocessing is used for computing the Mean Reciprocal Rank (MRR). This dataset is commonly used in previous Code Search studies [12, 17].

5.3 Evaluation Metrics

The Code Summarization task is evaluated using BLEU [34] and METEOR [5] where the generated summaries are compared against the ground truth comments. These metrics are commonly used in

Code Summarization studies [43, 49]. BLEU and METEOR scores are numbers between 0 and 1, and we report their percentages following previous work [28, 49].

BLEU is a precision-based metric and measures the n-gram geometric precision between the generated summary and the ground truth summary [35].

METEOR measures the alignment between the generated summary and the ground-truth summary using exact, stem, and synonym matches between words and phrases [6].

The Code Search task is a search task and it is evaluated using the Mean Reciprocal Rank (MRR).

MRR evaluates the Code Search task where it produces a list of possible responses to the code query. The reciprocal rank of the code query response is the multiplicative inverse of the rank of the first correct answer, and the MRR is the average of the reciprocal ranks of results for the code queries [38, 42].

5.4 Qualitative Analysis for Code Summarization

For Code Summarization, we further conducted a qualitative analysis. We randomly select 800 generated summaries (for each PLM) along with their original code, 100 pairs for each of the best performing monolingual and multilingual PLMs, following prior research [14, 29]. Amazon Mechanical Turk (MTurk) workers were hired to rate the quality of the generated summaries. The MTurkers rated the summary voluntarily, and for each rated summary, the MTurkers are given a compensation of one cent. There are three different annotators for each generated code summary, and different generated code summaries may not get the same annotators, due to the randomness of the MTurkers' assignments. We also ask the MTurkers if they understand the code and we only accept those ratings when the MTurkers have stated that they do understand the code. This will ensure that the annotations are proper and they will not be biased towards specified annotators. In total, 2,088 MTurkers have participated in the study. We used four common criteria to evaluate the summarization quality [29]:

Informativeness How well the summary capture the key points of the code?

Relevance Are the summary details consistent with in the code?

Fluency Are the summaries well-written and grammatically correct?

Comprehension Can the summaries help in understanding the code?

Three different workers were required to rate each summary between one and five, where one is the worst and five is the best.

6 RESULTS

In this section, we will first discuss the result of Code Summarization before Code Search.

6.1 RQ1: Does training and fine-tuning on the individual programming languages help to improve the performance over multilingual PLMs that are fine-tuned on multilingual datasets?

6.1.1 Code Summarization. Table 3 shows the best performing monolingual PLM. The first column shows the PLM and the second to sixth column show the BLEU and the METEOR metrics.

We observe that for all the PLMs, fine-tuning on the combined dataset gives the best performance. Comparing the monolingual and multilingual PLMs i.e., CodeBERT and GraphCodeBERT, the latter has better performance.

Table 4 shows the annotation results from the Amazon MTurkers. Generally, the MTurkers find that the generated summaries of all the PLMs are informative, relevant, fluent, and that they can help in the understanding of the code – the summaries are rated above 4 in these areas.

6.1.2 Code Search. Table 5 shows the Code Search results of the best performing fine-tuned model of each PLM. The first and second column show the model and its MRR scores. The third and fourth column show the improvement over CodeBERT and GraphCodeBERT respectively. Similar to the Code Summarization task, we observe that for all the PLMs, fine-tuning on the combined multilingual dataset give the best performance. However, we were surprised to observe that all the monolingual PLMs outperformed CodeBERT in MRR between 5% and 35.1%, and GraphCodeBERT in MRR between 2.3% and 32.6%.

We conducted a Mann Whitney U-test between using the monolingual datasets and the multilingual datasets for fine-tuning, and the p-value is $0.00256 < 0.05$, showing that our experiments are statistically significant.

Multilingual PLMs fine-tuned on the combined multilingual dataset perform better than monolingual PLMs in Code Summarization whereas monolingual PLMs perform better than multilingual PLMs on Code Search.

6.2 RQ2: Which PLM has the best Performance-to-Time Ratio?

6.2.1 Code Summarization. Figure 1 shows the Performance-to-Time Ratio of the PLMs: the ratio of BLEU-4 and the model fine-tuning time, and the ratio of METEOR and the model fine-tuning time. We observe that for the multilingual PLMs, they have lower Performance-to-Time Ratio than the other PLMs. This shows that although the multilingual PLMs have higher performance, it would take much longer to fine-tune them. For the monolingual PLMs, we observe that many of them have much larger Performance-to-Time Ratio than the multilingual PLMs, and monolingual PLM trained on the PHP dataset has the highest scores.

6.2.2 Code Search. Figure 2 shows the Performance-to-Time Ratio of MRR and the model fine-tuning time. Similarly, the MRR and the model fine-tuning time are normalized to a value between 0 and 1 before the ratios are computed. GraphCodeBERT and CodeBERT refer to the multilingual PLMs while the other PLMs refer to the monolingual PLMs. Similar to Code Summarization, we observe that the monolingual PLMs have higher Performance-to-Time Ratio than the multilingual PLMs. For the monolingual PLMs, we observed that the monolingual PLM trained on the Ruby dataset has the highest score.

Table 3: Code Summarization using different PLMs (trained using monolingual dataset) fine-tuned on the monolingual and combined multilingual datasets for Ruby. Here, we only show the fine-tuned model of each PLM having the best BLEU scores e.g., $PLM_{Ruby}/Combined$ refers to the PLM pre-trained using the Ruby dataset and fine-tuned on the combined multilingual dataset. The left value in the bracket shows the average score while the right value shows the standard deviation.

PLM	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR
$PLM_{Ruby}/Combined$	16.5 (15.0/2.3)	7.9 (6.3/1.4)	4.3 (3.1/0.9)	2.6 (1.7/0.6)	12.5 (11.0/0.1)
$PLM_{JavaScript}/Combined$	15.7 (14.6/1.7)	7.4 (6.2/1.5)	4.0 (3.2/1.1)	2.4 (1.8/0.1)	12.4 (11.2/1.1)
$PLM_{PHP}/Combined$	16.1 (14.2/2.7)	7.9 (6.1/1.6)	4.5 (3.1/1.0)	2.8 (1.8/0.7)	12.8 (11.5/0.9)
$PLM_{Java}/Combined$	16.1 (14.1/2.4)	8.0 (6.3/1.6)	4.6 (3.3/1.1)	2.8 (1.9/0.8)	12.7 (11.5/1.0)
$PLM_{Python}/Combined$	15.7 (15.1/2.2)	7.5 (6.7/1.3)	4.2 (3.6/0.9)	2.6 (2.0/0.6)	12.5 (11.8/0.7)
$PLM_{Go}/Combined$	14.7 (15.0/2.7)	7.4 (6.4/1.6)	4.2 (3.3/1.1)	2.6 (1.8/0.7)	12.6 (11.3/1.0)
CodeBERT/Combined	15.4	8.2	4.9	3.1	13.2
GraphCodeBERT/Combined	16.5	9.1	5.8	3.9	13.5




Figure 1: Performance-to-Time Ratio (PTR) for Code Summarization. Monolingual PLMs have higher BLEU-4 and METEOR PTR.




Figure 2: PTR for Code Search. The monolingual PLMs have shown a higher MMR Performance-to-Time Ratio (PTR).

Table 4: Qualitative results from the MTurker studies.

	Info.	Rel.	Flu.	Compre.
$PLM_{Ruby}/Combined$	4.40	4.57	4.63	4.62
$PLM_{JavaScript}/Combined$	4.46	4.47	4.54	4.57
$PLM_{PHP}/Combined$	4.49	4.45	4.53	4.58
$PLM_{Java}/Combined$	4.42	4.40	4.44	4.45
$PLM_{Python}/Combined$	4.47	4.42	4.53	4.63
$PLM_{Go}/Combined$	4.45	4.45	4.42	4.49
CodeBERT/Combined	4.41	4.46	4.46	4.41
GraphCodeBERT/Combined	4.33	4.35	4.44	4.39

Table 5: Code Search using monolingual PLMs fine-tuned on the monolingual and combined multilingual dataset. Here, we only show the best performing fine-tuned model of each PLM. The left value in the bracket shows the average score while the right value shows the standard deviation.

PLM	MRR	Improve CodeBERT	Improve GCodeBERT
$PLM_{Ruby}/Combined$	0.57(0.41/0.12)	+35.1%	+32.6%
$PLM_{JavaScript}/Combined$	0.44(0.32/0.07)	+5%	+2.3%
$PLM_{PHP}/Combined$	0.47(0.29/0.09)	+9.9%	+9.3%
$PLM_{Java}/Combined$	0.46(0.29/0.10)	+9.8%	+7%
$PLM_{Python}/Combined$	0.48(0.31/0.10)	+12.7%	+11.6%
$PLM_{Go}/Combined$	0.46(0.24/0.13)	+9.8%	+7%

We conducted a Mann Whitney U-test between using the monolingual datasets and the multilingual datasets for the Performance-to-Time Ratio, and the p-value is $0.00256 < 0.05$, showing that our experiments are statistically significant.

Monolingual PLMs fine-tuned on the combined dataset have the best Performance-to-Time Ratio. Researchers should consider choosing monolingual PLMs if they were to pre-train a PLM from scratch.

6.3 RQ3: What are the best settings for zero-shot downstream tasks?

Table 6: Zero-shot Code Summarization using different monolingual PLMs (other than Ruby) fine-tuned on the monolingual dataset (other than Ruby). Here, we only show the best performing fine-tuned model of each PLM (the average and standard deviation scores of other fine-tuned models are shown as the left and right values within the braces) e.g., $PLM_{JavaScript/Python}$ refers to the PLM pre-trained using the Javascript dataset and fine-tuned on the Python dataset.

PLM	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR
$PLM_{JavaScript/Python}$	14.8 (13.8/0.9)	7.3 (5.7/1.3)	4.0 (2.8/1.0)	2.4 (1.5/0.7)	12.3 (10.9/1.2/)
$PLM_{PHP/Python}$	16.2 (13.5/3.0)	7.6 (5.6/1.6)	4.0 (2.8/0.9)	2.4 (1.5/0.6)	12.5 (11.3/0.8)
$PLM_{Java/Python}$	16.2 (13.5/2.4)	7.9 (5.7/1.5)	4.5 (2.9/1.0)	2.8 (1.7/0.7)	12.6 (11.2/0.1)
$PLM_{Python/Python}$	15.7 (14.3/1.9)	7.6 (6.2/1.2)	4.2 (3.2/0.8)	2.6 (1.8/0.5)	12.5 (11.6/0.8)
$PLM_{Go/Python}$	16.4 (14.3/2.5)	7.9 (5.9/1.5)	4.3 (2.9/1.0)	2.5 (1.5/0.7)	12.4 (11.0/0.9)

6.3.1 Code Summarization. Table 6 shows the zero-shot Code Summarization using the monolingual PLMs. The first column shows the PLM while column two to six show the BLEU and METEOR scores respectively. For all the PLMs, we observe that PLMs that are fine-tuned on the Python dataset has the best performance over other monolingual datasets. The average and standard deviation of the other fine-tuned models for each PLM is shown on the left and right values inside the braces.

Table 7: Zero-Shot Code Search using different monolingual PLMs (other than Ruby) fine-tuned on the monolingual dataset (other than Ruby). Here, we only show the best performing fine-tuned model of every PLM.

PLM	MRR
$PLM_{JavaScript/Python}$	0.37 (0.29/0.05)
$PLM_{PHP/Python}$	0.31 (0.25/0.05)
$PLM_{Java/Python}$	0.34 (0.24/0.07)
$PLM_{Python/Python}$	0.36 (0.26/0.08)
$PLM_{Go/Python}$	0.36 (0.19/0.11)

6.3.2 Code Search. Table 7 shows the zero-shot results for Code Search using the monolingual PLMs. The first column shows the PLM while the second column shows the MRR scores. Interestingly, we observed that similar to the zero-shot settings in Code Summarization, fine-tuning on the Python dataset has the best MRR performance.

We conducted a Mann Whitney U-test between using the datasets containing Ruby and the datasets excluding Ruby, and the p-value is $0.00604 < 0.05$, showing that our experiments are statistically significant.

For the zero-shot settings, we observed that PLMs fine-tuned on the Python dataset has the best performance. Researchers should consider using the Python dataset to fine-tune PLMs for Ruby.

6.4 RQ4: What effect does the PLMs have on different code length?

6.4.1 Code Summarization. Table 8 shows the effects of the different code lengths when the PLMs are used on Code Summarization. For code lengths within the first percentile and code lengths above

the third percentile, we observe that majority of the PLMs have the best performance when fine-tuned on the combined multilingual datasets. However, we observed that for the other code lengths, some of the PLMs perform better when fine-tuned on Ruby or Python dataset.

6.4.2 Code Search. Table 9 shows the effects of the different code lengths when the PLMs are used in Code Search. We observed that for all the different PLMs, fine-tuning on the combined multilingual dataset gave the best performance in all code lengths. Among the PLMs, there is very little variation in the MRR scores on the different code lengths i.e., their scores differ within +/-0.05 and majority of them are within +/-0.03.

For different code lengths, we observed similar performance when the PLMs are tested on all the code lengths.

6.5 RQ5: How effective is our strategy to decide in advance a language that can work well for a target low-resource language?

6.5.1 Code Summarization. Based on the suitability equation 1, Python, PHP, Java, Javascript and Go have scores of 0.55, 0.92, 0.14, 0.52 and 0.46, when compared to Ruby, respectively. Python, PHP and Go are over 0.5 and thus we selected them for fine-tuning. Ruby is also included as it is the target language. Table 10 shows the performance of BLEU and METEOR when the PLMs are fine-tuned on the proposed set of programming languages. We observe improvement in BLEU and METEOR on both monolingual and multilingual PLMs. For BLEU-1, BLEU-2, BLEU-3, BLEU-4 and METEOR, the improvement ranges from 0.6% to 4.3%, 1.1% to 13.5%, 11.9% to 30%, 3.8% to 41.7%, and 1.5% to 5.6%, respectively.

6.5.2 Code Search. From Table 5, we observed that the PLMs fine-tuned on the combined multilingual dataset has the best performance in MRR. Specifically, we observed that all the monolingual PLMs have better performance in MRR than the multilingual PLMs. The improvement over CodeBERT and GraphCodeBERT ranges from 5% to 35.1% and 2.3% to 32.6%, respectively.

Table 8: Effects of PLMs on different code lengths fine-tuned on monolingual and combined multilingual dataset. Here, for each monolingual PLM, we show only the fine-tuned model of each PLM having the best BLEU scores e.g., $PLM_{JavaScript}/Combined$ refers to the PLM pre-trained using the Javascript dataset and fine-tuned on the combined multilingual dataset.

PLM	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR
0 - 1st Quartile Code Length					
$PLM_{Ruby}/Combined$	18.2	8.5	4.7	3.0	13.2
$PLM_{JavaScript}/Python$	16.7	8.2	4.6	2.9	12.9
$PLM_{PHP}/Combined$	18.3	9.1	5.4	3.6	13.8
$PLM_{Java}/Combined$	18.8	9.2	5.2	3.3	13.7
$PLM_{Python}/Ruby$	21.0	10.2	5.8	3.3	13.0
$PLM_{Go}/Combined$	17.7	8.8	5.1	3.4	13.7
CodeBERT/Combined	17.8	9.4	5.5	3.4	14.0
GraphCodeBERT/Combined	17.8	9.4	5.5	3.4	14.0
1st - 2nd Percentile Code Length					
$PLM_{Ruby}/Python$	16.4	7.9	4.3	2.6	12.7
$PLM_{JavaScript}/Ruby$	18.5	8.9	5.2	3.3	12.2
$PLM_{PHP}/Combined$	17.2	8.5	4.9	3.1	13.2
$PLM_{Java}/Combined$	18.0	9.3	5.5	3.5	13.2
$PLM_{Python}/Python$	17.4	8.7	5.2	3.4	13.2
$PLM_{Go}/Ruby$	19.8	9.1	5.1	2.9	11.7
CodeBERT/Combined	16.6	9.0	5.6	3.7	13.9
GraphCodeBERT/Combined	17.2	9.4	6.0	4.1	14.0
2nd - 3rd Quartile Code Length					
$PLM_{Ruby}/Combined$	14.2	7.0	4.1	2.6	11.7
$PLM_{JavaScript}/Ruby$	15.4	6.9	3.9	2.3	10.1
$PLM_{PHP}/Python$	14.1	7.0	3.9	2.4	12.0
$PLM_{Java}/Python$	13.5	6.9	4.2	2.7	12.0
$PLM_{Python}/Ruby$	16.3	7.7	4.3	2.5	11.5
$PLM_{Go}/Python$	14.2	7.0	3.9	2.3	11.8
CodeBERT/Combined	16.4	7.0	3.7	2.0	11.2
GraphCodeBERT/Combined	13.7	7.7	4.9	3.4	12.9
3rd Quartile - Max Code Length					
$PLM_{Ruby}/Combined$	16.2	7.5	4.1	2.2	11.9
$PLM_{JavaScript}/Python$	14.4	6.7	3.5	2.0	11.7
$PLM_{PHP}/Combined$	15.1	7.1	3.9	2.2	11.9
$PLM_{Java}/Combined$	14.5	6.7	3.5	1.9	11.7
$PLM_{Python}/Combined$	14.6	6.9	4.0	2.5	11.6
$PLM_G/Python$	15.1	7.3	4.2	2.5	11.4
CodeBERT/Combined	14.3	7.4	4.3	2.7	12.2
GraphCodeBERT/Combined	14.7	7.8	4.9	3.2	12.3

Our proposed strategies in selecting PLs for fine-tuning downstream tasks are effective. For Code Summarization, the improvement ranges from 0.6% to 41.7% in BLEU, and from 1.5% to 5.6% in METEOR. For Code Search, the MRR scores improve from 2.3% to 35.1%.

7 DISCUSSIONS

Which PLM should I use? In this work, we studied the performance impact of monolingual and multilingual PLMs. We found that although multilingual PLMs show good performance in the automatic metrics, the MTurkers do not find any major difference between the summaries generated from monolingual and multilingual PLMs for Code Summarization in the qualitative study. Furthermore, the Performance-to-Time Ratio of the PLMs also suggest that monolingual PLMs are more efficient. In the case of Code

Search, we observed that monolingual PLMs have better performance than multilingual PLMs. Thus, considering both efficiency and performance, we propose that developers do not merely use the multilingual PLMs in their task, but to also compare with the monolingual PLMs.

Different strategy in selecting a suitable PL for different downstream tasks In our experiment, we have observed that the strategies in selecting a suitable programming language for Code Summarization and Code Search works well in their respective tasks. However, we also observed that the strategy for Code Summarization may not work as well in Code Search, and vice-versa. We believe that depending on the task, a different strategy may be required. Nonetheless, we showed that our proposed strategies (Section 4) are effective and we still recommend developers to adopt our proposed strategies in any of their tasks before attempting to come out with a new one.

Table 9: Effects of PLMs on different code lengths fine-tuned on the monolingual dataset. Here, we only show the best performing fine-tuned model of each PLM e.g., $PLM_{JavaScript/Python}$ refers to the PLM pre-trained using the Javascript dataset and fine-tuned on the Python dataset.

0 - 1st Quartile		1st - 2nd Quartile		2nd - 3rd Quartile		3rd Quartile - Max	
PLM	MRR	PLM	MRR	PLM	MRR	PLM	MRR
$PLM_{Ruby}/Comb.$	0.57	$PLM_{Ruby}/Comb.$	0.59	$PLM_{Ruby}/Comb.$	0.58	$PLM_{Ruby}/Comb.$	0.55
$PLM_{JavaScript}/Comb.$	0.44	$PLM_{JavaScript}/Comb.$	0.47	$PLM_{JavaScript}/Comb.$	0.44	$PLM_{JavaScript}/Comb.$	0.43
$PLM_{PHP}/Comb.$	0.47	$PLM_{PHP}/Comb.$	0.47	$PLM_{PHP}/Comb.$	0.48	$PLM_{PHP}/Comb.$	0.44
$PLM_{Java}/Comb.$	0.35	$PLM_{Java}/Comb.$	0.34	$PLM_{Java}/Comb.$	0.32	$PLM_{Java}/Comb.$	0.33
$PLM_{Python}/Comb.$	0.48	$PLM_{Python}/Comb.$	0.48	$PLM_{Python}/Comb.$	0.47	$PLM_{Python}/Comb.$	0.48
$PLM_{Go}/Comb.$	0.47	$PLM_{Go}/Comb.$	0.48	$PLM_{Go}/Comb.$	0.45	$PLM_{Go}/Comb.$	0.43
CodeBERT/Comb.	0.44	CodeBERT/Comb.	0.42	CodeBERT/Comb.	0.42	CodeBERT/Comb.	0.42
GCodeBERT/Comb.	0.43	GCodeBERT/Comb.	0.44	GCodeBERT/Comb.	0.46	GCodeBERT/Comb.	0.41

Table 10: Code Summarization using different PLMs fine-tuned on a subset of selected languages (*Ruby + Python + PHP + Go*) and tested on Ruby. The value inside the braces symbolized the percentage improvement over PLMs fine-tuned on the combined multilingual dataset. There is an improvement in all the metrics among the different PLMs.

PLM	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR
PLM_{Ruby}	15.7 (-4.8%)	7.7 (-2.5%)	4.3 (+0%)	2.7 (+3.8%)	12.7 (+2.7%)
$PLM_{JavaScript}$	15.6 (+0.6%)	8.4 (+13.5%)	5.2 (+30%)	3.4 (+41.7%)	13.1 (+5.6%)
PLM_{PHP}	16.8 (+4.3%)	8.6 (+8.9%)	5.1 (+13.3%)	3.2 (+14.3%)	12.8 (+0%)
PLM_{Java}	16.6 (+3.1%)	8.8 (+10%)	5.4 (+17.4%)	3.5 (+25%)	13.1 (+3.1%)
PLM_{Python}	15.9 (+1.3%)	8.3 (+10.7%)	5.1 (+21.4%)	3.3 (+26.9%)	13.1 (+4.8%)
PLM_{Go}	14.8 (+0.7%)	7.8 (+5.4%)	4.7 (+11.9%)	3.2 (+23.1%)	12.8 (+1.6%)
CodeBERT	15.8 (+2.6%)	8.7 (+6.1%)	5.4 (+12.9%)	3.5 (+21.5%)	13.4 (+1.5%)
GraphCodeBERT	16.6 (+0.6%)	9.2 (+1.1%)	5.8 (+0%)	3.9 (+0%)	13.5 (+0%)

Non-exact Code Duplicates in CodeSearchNet Dataset We note that non-exact code duplicates (Type 2-4 code clones) may exist in the CodeSearchNet dataset. We believe that excluding them during pre-training and fine-tuning may not necessarily yield more robust models, and that the PLMs can benefit from learning more diverse code structure using the dataset that contains non-exact code duplicates.

8 THREATS TO VALIDITY

External Validity. In this study, we discuss the results of different settings for Code Summarization and Code Search in Go, Java, Javascript, PHP, Python and Ruby. The tasks and the programming languages in our study are restricted, and the results might not be generalizable to other programming languages and tasks.

Internal Validity. We process the publicly available datasets, following other research [12, 22]. A potential threat may be related to not reaching the optimal performance of the pre-trained models, thus, having an under-trained PLM. We note that in the literature, there is no hard rule on the number of training steps for pre-training and to determine an optimal stopping criteria is still an open problem. Existing studies use fixed number of steps (a fraction of the dataset) as a stopping criteria [11, 12, 30]. For consistency, we pre-train all the PLMs involving different datasets for 50 epochs.

Construct Validity. The validity threat here can be related to the measures used to evaluate the results. To mitigate the bias that

might be related to a specific evaluation metric, we used multiple metrics that are commonly used in the downstream tasks.

9 CONCLUSION AND FUTURE WORKS

In this work, we studied both monolingual and multilingual PLMs on Code Summarization and Code Search. We observed that the monolingual PLMs have better Performance-to-Time Ratio, as compared to the multilingual PLMs. In addition, our proposed strategies in selecting a suitable programming language for Code Summarization and Code Search are efficient and can improve the current state-of-the-art performance. Based on these findings, we suggest the following: 1) to consider using a monolingual PLM fine-tuned on a combined multilingual dataset as it has a higher Performance-to-Time Ratio, and 2) use our proposed strategies for Code Summarization and Code Search. In future, we aim to study the applicability of PLMs on more downstream tasks, and generalizing a strategy that can work well in multiple downstream tasks.

ACKNOWLEDGMENTS

This research is support by a grant from Natural Sciences and Engineering Research Council of Canada RGPIN-2019-05175.

REFERENCES

- [1] 2021. Stack Overflow Developer Survey 2021. <https://insights.stackoverflow.com/survey/2021>
- [2] 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*.
- [3] Toufique Ahmed and Premkumar Devanbu. 2021. Multilingual training for Software Engineering. *arXiv preprint arXiv:2112.02043* (2021).
- [4] Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. 2021. Program Synthesis with Large Language Models. *arXiv preprint arXiv:2108.07732* (2021).
- [5] Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*.
- [6] Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*.
- [7] Nghi DQ Bui, Yijun Yu, and Lingxiao Jiang. 2021. Self-Supervised Contrastive Learning for Code Retrieval and Summarization via Semantic-Preserving Transformations. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [8] Matteo Ciniselli, Nathan Cooper, Luca Pascarella, Antonio Mastropaolo, Emad Aghajani, Denys Poshyvanyk, Massimiliano Di Penta, and Gabriele Bavota. 2021. An Empirical Study on the Usage of Transformer Models for Code Completion. *IEEE Transactions on Software Engineering* (2021).
- [9] Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations*.
- [10] Colin B Clement, Dawn Drain, Jonathan Timcheck, Alexey Svyatkovskiy, and Neel Sundaresan. 2020. PyMT5: multi-mode translation of natural language and Python code with transformers. (2020).
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. [arXiv:1810.04805](https://arxiv.org/abs/1810.04805) [cs.CL]
- [12] Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, et al. 2020. Codebert: A pre-trained model for programming and natural languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, Online.
- [13] David Gros, Hariharan Sezhiyan, Prem Devanbu, and Zhou Yu. 2020. Code to Comment “Translation”: Data, Metrics, Baseline & Evaluation. In *2020 35th IEEE/ACM International Conference on Automated Software Engineering*.
- [14] Max Grusky, Mor Naaman, and Yoav Artzi. 2018. Newsroom: A Dataset of 1.3 Million Summaries with Diverse Extractive Strategies. *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)* (2018).
- [15] Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, Shujie Liu, Long Zhou, Nan Duan, Alexey Svyatkovskiy, Shengyu Fu, et al. 2021. Graphcodebert: Pre-training code representations with data flow. In *International Conference on Learning Representations*.
- [16] Prakhar Gupta, Matteo Pagliardini, and Martin Jaggi. 2019. Better Word Embeddings by Disentangling Contextual n-Gram Information. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*.
- [17] Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. 2019. Codesearchnet challenge: Evaluating the state of semantic code search. *arXiv preprint arXiv:1909.09436* (2019).
- [18] Xue Jiang, Zhuoran Zheng, Chen Lyu, Liang Li, and Lei Lyu. 2021. TreeBERT: A Tree-Based Pre-Trained Model for Programming Language. In *Uncertainty in Artificial Intelligence*.
- [19] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. Tinybert: Distilling bert for natural language understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2020*.
- [20] Toshihiro Kamiya, Shinji Kusumoto, and Katsuro Inoue. 2002. CCFinder: A multilinguistic token-based code clone detection system for large scale source code. *IEEE Transactions on Software Engineering* (2002).
- [21] Aditya Kanade, Petros Maniatis, Gogul Balakrishnan, and Kensen Shi. 2019. Pre-trained contextual embedding of source code. (2019).
- [22] Aditya Kanade, Petros Maniatis, Gogul Balakrishnan, and Kensen Shi. 2020. Learning and evaluating contextual embedding of source code. In *International Conference on Machine Learning*.
- [23] Rafael-Michael Karampatsis and Charles Sutton. 2020. Scelmo: Source code embeddings from language models. *arXiv preprint arXiv:2004.13214* (2020).
- [24] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves 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* (2020).
- [25] Jia Li, Yongmin Li, Ge Li, Xing Hu, Xin Xia, and Zhi Jin. 2021. EDITSUM: A Retrieve-and-Edit Framework for Source Code Summarization. In *2021 36th IEEE/ACM International Conference on Automated Software Engineering*.
- [26] Liuqing Li, He Feng, Wenjie Zhuang, Na Meng, and Barbara Ryder. 2017. Ccleaner: A deep learning-based clone detection approach. In *2017 IEEE International Conference on Software Maintenance and Evolution*.
- [27] Fang Liu, Ge Li, Yunfei Zhao, and Zhi Jin. 2020. Multi-task learning based pre-trained language model for code completion. In *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering*.
- [28] Shangqing Liu, Yu Chen, Xiaofei Xie, Jing Kai Siow, and Yang Liu. 2021. Retrieval-Augmented Generation for Code Summarization via Hybrid GNN. In *International Conference on Learning Representations*.
- [29] Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- [30] 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:1907.11692](https://arxiv.org/abs/1907.11692) [cs.CL]
- [31] Junayed Mahmud, Fahim Faisal, Raihan Islam Arnob, Antonios Anastasopoulos, and Kevin Moran. 2021. Code to Comment Translation: A Comparative Study on Model Effectiveness & Errors. *The Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP 2021)* (2021).
- [32] Antonio Mastropaolo, Simone Scalabrino, Nathan Cooper, David Nader Palacio, Denys Poshyvanyk, Rocco Oliveto, and Gabriele Bavota. 2021. Studying the usage of text-to-text transfer transformer to support code-related tasks. In *2021 IEEE/ACM 43rd International Conference on Software Engineering*.
- [33] Pengyu Nie, Jiyang Zhang, Junyi Jessy Li, Raymond J Mooney, and Milos Gligoric. 2021. Evaluation Methodologies for Code Learning Tasks. *arXiv preprint arXiv:2108.09619* (2021).
- [34] 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*. Association for Computational Linguistics, Philadelphia, Pennsylvania, USA, 311–318.
- [35] 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*.
- [36] Michael Pradel and Koushik Sen. 2018. Deepbugs: A learning approach to name-based bug detection. *Proceedings of the ACM on Programming Languages OOPSLA* (2018).
- [37] Xipeng Qiu, Tianxiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. 2020. Pre-trained models for natural language processing: A survey. *Science China Technological Sciences* (2020), 1–26.
- [38] Dragomir R Radev, Hong Qi, Harris Wu, and Weiguo Fan. 2002. Evaluating Web-based Question Answering Systems.. In *Proceedings of the Third International Conference on Language Resources and Evaluation*.
- [39] 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. *Journal of Machine Learning Research* (2020).
- [40] Baptiste Roziere, Marie-Anne Lachaux, Lowik Chanasusot, and Guillaume Lample. 2020. Unsupervised Translation of Programming Languages. *Advances in neural information processing systems* (2020).
- [41] Alexey Svyatkovskiy, Shao Kun Deng, Shengyu Fu, and Neel Sundaresan. 2020. Intellicode 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*.
- [42] Ellen M Voorhees et al. 1999. The TREC-8 question answering track report. In *Trec*.
- [43] 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*.
- [44] Xin Wang, Fei Mi Yasheng Wang, Pingyi Zhou, Yao Wan, Xiao Liu, Li Li, Hao Wu, Jin Liu, and Xin Jiang. 2022. SYNCOBERT: Syntax-Guided Multi-Modal Contrastive Pre-Training for Code Representation. (2022).
- [45] Yanlin Wang, Lun Du, Ensheng Shi, Yuxuan Hu, Shi Han, and Dongmei Zhang. 2020. *CoCoGUM: Contextual Code Summarization with Multi-Relational GNN on UMLs*. Technical Report MSR-TR-2020-16. Microsoft.
- [46] 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, EMNLP 2021.*
- [47] 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*.
- [48] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems* (2019).
- [49] Jian Zhang, Xu Wang, Hongyu Zhang, Hailong Sun, and Xudong Liu. 2020. Retrieval-Based Neural Source Code Summarization. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*. 13 pages.