Multilingual training for Software Engineering

Toufique Ahmed
University of California, Davis
Davis, California, USA
tfahmed@ucdavis.edu

Premkumar Devanbu
University of California, Davis
Davis, California, USA
ptdevanbu@ucdavis.edu

ABSTRACT

Well-trained machine-learning models, which leverage large amounts of open-source software data, have now become an interesting approach to automating many software engineering tasks. Several SE tasks have all been subject to this approach, with performance gradually improving over the past several years with better models and training methods. More, and more diverse, *clean, labeled* data is better for training; but constructing good-quality datasets is time-consuming and challenging. Ways of augmenting the volume and diversity of clean, labeled data generally have wide applicability. For some languages (*e.g.,* Ruby) labeled data is less abundant; in others (*e.g.,* JavaScript) the available data maybe more focused on some application domains, and thus less diverse. As a way around such data bottlenecks, we present evidence suggesting that human-written code in different languages (which performs the same function), is rather similar, and particularly preserving of identifier naming patterns; we further present evidence suggesting that identifiers are a very important element of training data for software engineering tasks. We leverage this rather fortuitous phenomenon to find evidence that available *multilingual* training data (across different languages) can be used to amplify performance. We study this for 3 different tasks: code summarization, code retrieval, and function naming. We note that this data-augmenting approach is broadly compatible with different tasks, languages, and machine-learning models.

CCS CONCEPTS

• **Software and its engineering** → **Software notations and tools**; • **Computing methodologies** → **Machine learning**.

KEYWORDS

code summarization, code search, method name prediction, deep learning

ACM Reference Format:
Toufique Ahmed and Premkumar Devanbu. 2022. Multilingual training for Software Engineering. In *44th International Conference on Software Engineering (ICSE '22), May 21–29, 2022, Pittsburgh, PA, USA.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3510003.3510049

ICSE '22, May 21–29, 2022, Pittsburgh, PA, USA

ACM ISBN 978-1-4503-9221-1/22/05.
https://doi.org/10.1145/3510003.3510049

1 INTRODUCTION

Researchers in the NLP area have reported that multilingual training is beneficial for low-resource language [16, 23, 57, 63]. Several papers show that multilingual-trained models show better performance [36, 62] and are more practical to deploy [9]. However, this is observed in two situations: 1) for low-resource languages and 2) when the languages are related. We find that programs in different languages solving the same problem use more similar identifiers; furthermore different languages sometimes have similar keywords and operators. High capacity deep learning models are capable of learning *interlingua*: shared semantic representation between languages [34]. Moreover, with tasks like summarization, or method naming, we are dealing with a simplified, many-to-one setting: translating multiple source languages to a single target language), which is believed to be easier than multi-way task [20, 76]. We begin by introducing the code summarization task, which we use to motivate multilingual training.

Developers often rely heavily on comments, to gain a quick (even if approximate) understanding of the specification and design of code they are working on. An actual example of a comment is shown in Figure 1. Such comments help a developer gain a quick mental preview of *what* the proximate code does, and *how* it might go about it; this helps the developer know what to look for in the code. Knowing that such comments are useful to others (or even later to oneself) incentivizes developers to create comments that explain the code; however the resulting redundancy (*viz.,* code that does something, and some nearby English text that describes just what the code does), with the same concept expressed in two languages results in a bit of extra work for the original coder.
This extra work, of creating aligned comments explaining the code, can be fruitfully viewed [21] as a task related to *natural language translation* (NLT) (*e.g.,* translating English to German). The mature & powerful technology of NLT becomes applicable for comment synthesis; ML approaches developed for the former can be used for the latter. An effective comment synthesizer could help developers: by saving them the trouble of writing comments; and perhaps even be used on-demand in the IDE to create descriptions of selected bits of code.

Comment synthesis is now an active research area, including many projects such as CodeNN [30], DeepCom [26], Astattgru [40], CodeBERT [18], Rencos [74], SecNN [42], PLBART [1], CoTexT [54], ProphetNet-X [55], NCS [2], Code2seq [7], Re^{2}Com [71], and many more [19, 24, 25, 27, 28, 38, 39, 41, 49, 50, 66, 67, 69, 70, 72, 73]. All these approaches rely on datasets of aligned code-comment pairs. Typically, these datasets are then used to train complex deep learning models to model a probabilistic distribution of the form *p(comments | code)* ; one can sample from these (usually generative) models to create candidate comments for a given a piece of code. Given a dataset of code-comment pairs in a specific language, *e.g.,*

Java, or Python, or PHP, or Ruby, one can train models to translate code in *that* language to comments. The quality of the translation will depend largely upon the inductive power of the model, and quality and diversity of the code-comment dataset.

```
//Returns the text content of
//this node and its descendants.
public String getTextContent() {
    StringBuilder sb=new StringBuilder(getChildNodesCount()+1);
    appendTextContent(sb);
    return sb.toString();
}
```

Figure 1: ***Example for code comment generation task***

Of late, given the power of GPUs, and the capacity of the models, the limitations largely arise from dataset quality and diversity, especially in languages for which limited, or rather specialized data is available. For instance, CodeXGLUE [47] dataset consists of six languages (*i.e.*, Ruby, Java, JavaScript, Go, Php, Python). Most languages have well over 100,000 training examples, covering a wide set of application domains. Some languages, particularly Ruby and Javascript, have far fewer examples, and cover a narrower range of application domains. As a result, state-of-the-art models perform less well for these two languages. This is a well-known problem for natural language translation: while training data for language pairs like $English \leftrightarrow French$ is abundant, resources may be lacking for less-used languages like *Quechua* or *Badaga*. In such cases, a common technique is adapt ML models to learn useful statistics from abundant data in other, perhaps related languages [51]. This works well when languages often have similar grammars, and share common word etymologies.

We propose an analogous approach to improve the diversity and quality of training data for software-engineering tasks, exploiting an interesting property of source code that human beings write. It's generally agreed that variable names help code comprehension [37]. Developers know this, and typically choose descriptive variable names (reflective of code logic and purpose) regardless of the language they are coding in. Thus, one could expect that developers coding *the same functionality, using similar algorithms, even in different languages, will use similar variable names.* This suggests that machine-learning approaches could sometimes leverage corpora in different programming languages. This paper a) shows that this expectation actually has a sound empirical basis, and then b) demonstrates that this approach in fact works not just for code summarization, but also for several other tasks. We make the following contributions.

(1) Using the RosettaCode dataset, we provide evidence that programs solving the same problem in different languages are more likely to use the same or similar identifier names.
(2) We show evidence suggesting that cross-language training (*e.g.,* train on Python, test on Ruby) can sometimes lead to better performance than same-language training.
(3) We study the relative value of identifiers and syntax, using ablation, and find that identifier names may matter more.
(4) We show that pooled multilingual training data improves performance on several tasks, but especially for languages lacking in diverse and abundant data. We top a leaderboard for code-comment synthesis[1].
(5) We show that multilingual training helps for two other tasks: *code retrieval*, and *method name prediction.*
(6) Finally, we evaluate a few different design choices for multilingual training, and discuss threats to our findings.

Overall, this paper a) shows that multilingual training is yet another useful technique in the general arsenal of ML approaches to exploit the naturalness of code, b) shows why it is useful, and c) shows how to take good advantage of it.

Note: Technical details follow, but precisely: what we study here is multilingual training in the *fine-tuning stage of "foundation models" [12].* Foundation models for code, like CodeBERT, GraphCodeBERT [18, 22] already use multilingual data for *pre-training.* While pre-training is self-supervised and is done with unlabeled corpora, task-specific fine-tuning is usually supervised, using clean, hard-won labeled data; multilingual pooling can be useful here.

2 BACKGROUND & MOTIVATION

We now present some motivating evidence suggesting the value of multilingual training data for deep-learning applications to software tasks. We begin the argument focused on code summarization.

Deep learning models have been widely applied to code summarization, with papers reporting substantial gains in performance over recent years [1, 2, 7, 18, 19, 24–28, 30, 38–42, 50, 54, 55, 66, 67, 69–74]. We focus here on *what information in the code* ML models leverage for summarization (while we use summarization to motivate the approach, we evaluate later on 3 different tasks). Does every token in the program under consideration matter, for the code summarization task? Or, are the function and variable names used in the programs most important? Since identifiers carry much information about the program, this may be a reasonable assumption.

Considering the content words[2] in the example in Figure 1 there are four major terms (*i.e.*, Returns, text content, node, and descendants) used in the summary. The first 3 directly occur as tokens or subtokens in the code. Though the word "descendants" is missing in the program, high capacity neural models like BERT [17] can learn to statistically connect, *e.g.,* "descendant" with the identifier subtoken "child". This suggests that, perhaps, comments are recoverable primarily from identifiers. If this is so, and identifiers matter more for comments than the exact syntax of the programming language, that may actually be very good news indeed. If developers choose identifiers in the same way across different languages (*viz.*, problem-dependent, rather than language dependent) perhaps we can improve the diversity and quality of dataset by pooling training set across may languages. Pooled data sets may allow us to fine-tune using multilingual data, and improve performance, especially for low-resource languages (*e.g.*, Ruby and JavaScript from

[1]This claim is based on publicly available evidence. Please check *https://microsoft.github.io/CodeXGLUE/*

[2]"Content" words in linguistics, are words that carry meaning, as contrasted with *function* words, such as prepositions, pronouns, and conjunctions, which denote grammatical relationships. See https://en.wikipedia.org/wiki/Content_word. In code, we consider function words to be keywords, operators and punctuations, and content words to be identifiers (functions, variables, types, *etc*)

CodeXGLUE [47]). Since this is a core theoretical background for our work, we start off with two basic research questions to empirically gauge the possibility and promise of multilingual fine-tuning.

RQ1 What role do identifiers play in for code summarization?
RQ2 Do programs that solve the same problem in different languages tend to use similar identifier names?

2.1 RQ1: Role played by identifiers

We first examine the importance of *identifiers* for code summarization; specifically, we compare the relative value of identifier tokens and other tokens. We use the CodeXGLUE dataset and pre-trained CodeBERT embeddings for the task [18]. We begin with a brief backgrounder on CodeBERT [18] & BERT [17].

CodeBERT uses the pre-training + fine-tuning strategy of BERT, RoBERTa *etc* [17, 45]. This approach begins with a self-supervised "pre-training" step, to learn textual patterns from a large, unlabeled, corpus using just the content; in the next step, "fine-tuning", task-specific *labeled* data is used to provide task-related supervised training. This approach is known to achieve state-of-the-art performance in both natural language processing, and software-related tasks [3, 4, 11, 18, 22, 32, 33, 35, 50, 75].

We study the effect of identifiers in several steps. For the pre-training step, we start with the available CodeBERT model, which is pre-trained on a large, multilingual corpus of code. For the fine-tuning step, for this task, we use the CodeXGLUE benchmark dataset (see table 4 for languages and dataset sizes); we start with the original set of code-comment pairs, and apply two different treatments to create overall three different fine-tuning training datasets–1) base case leaving code as is, 2) a treatment to emphasize identifiers, and 3) a treatment to de-emphasize them. First, to emphasize identifiers we abstract out the program's keywords, separators, and operators by replacing those with three generic tokens (*i.e.*, "key", "sep", and "opt"), thus forcing the model (during fine-tuning) to rely more on the identifiers, for the task. Next, to assess the importance of keywords, separators, and operators, we abstract out the identifiers with a generic token "id". We fine-tune the model separately after each of these abstraction steps, thus yielding 3 fine-tuned models: the baseline, keyword-abstracted, and identifier-abstracted. We compare the results (smoothed BLEU-4) across all three.

If a fine-tuned model's performance is relatively unaffected by an abstraction, one may infer that the model relies less on the abstracted tokens. We perform these experiments with two languages with low-resource (*i.e.*, Ruby and JavaScript, See table 4) and two languages with high-resource (*i.e.*, Java and Python). We train, validate, and test with the same dataset in each case. For each test instance, we have one value from the complete program and another one from each of the two abstracted versions. We compared these values, using two distinct pair-wise Wilcoxon tests: 1) Alternative Hypothesis (AH): complete program > identifier de-emphasis & 2) AH: complete program > identifier emphasis. We also perform the same test with the keyword-abstracted and identifier-abstracted versions (AH: identifier emphasis > identifier de-emphasis).

The data (table 1) suggests that abstracting the keyword, separator, and operator has a smaller impact on the performance: the BLEU-4 scores are rather similar (with *effect size* ranging from 0.002 to 0.033) to those from the unabstracted code. On the other hand, when de-emphasizing identifiers, the performance drops more, with effect sizes 5x-100x larger. We find similar results while comparing the emphasizing and de-emphasizing identifiers versions (omitted for brevity).

Dataset	Complete Program	Abstracting keyword, operator, separator			Abstracting identifiers		
	BLEU-4	BLEU-4	Effect Size	p-value (adjusted)	BLEU-4	Effect Size	p-value (adjusted)
Ruby	**12.53**	11.57	-0.028	0.008	7.94	-0.238	<0.001
JavaScript	**13.86**	13.06	-0.033	<0.001	9.06	-0.175	<0.001
Java	**18.72**	18.72	-0.002	0.344	11.41	-0.254	0
Python	**18.25**	18.10	-0.010	<0.001	11.68	-0.288	0

Table 1: ***Role played by identifiers***

Language		Training					
		Ruby	JavaScript	Java	Go	PHP	Python
Testing	Ruby	*12.53*	11.84	**13.42**	12.32	**13.84**	**14.09**
	JavaScript	11.98	*13.86*	**14.16**	12.55	**13.90**	**14.09**
	Java	13.38	14.57	***18.72***	14.20	16.27	16.20
	Go	11.68	11.24	13.61	***18.15***	12.70	13.53
	PHP	17.52	19.95	22.11	18.67	***25.48***	21.65
	Python	14.10	14.44	16.77	14.92	16.41	***18.25***

Table 2: ***Intra and inter language training and testing***

The results in table 1 suggests that syntax is less relevant that identifier names. In all the prior works, the training and testing were done in the same language. Since syntax is less important, could we train and test with different languages? The CodeXGLUE dataset enables just such an experiment. Using six different languages, we apply a CodeBERT model fine-tuned in each language, to a test set in another language. Table 2 shows that for high-resource languages (*i.e.*, Java, go, PHP, and Python), we achieve the best result (diagonal) when training and test data are from the same language. However, the performance does not degrade to a very large extent when trained with one language and tested on a different one. Surprisingly we observe that for Ruby and JavaScript, we actually *achieve higher performance while trained with Java, PHP, and Python than the language itself.* That indicates that code summarization is not completely dependent on syntax (perhaps it relies more on identifier similarity, which we shall explore next)

Finding 1. Code summarization sometimes appears to train quite well with data sets from other languages, even if the syntax is different.

2.2 RQ2: Identifier similarity across Languages

Here, we evaluate RQ2: given a problem, do developers choose similar, descriptive identifiers, regardless of the programming language? Based on the findings in the previous section: if identifiers were indeed used in similar ways, perhaps code-comment pairs *from any programming language* could help train a code summarization model, *for any other language*. As an example, Figure 2 presents that all the "indexOf" functions implemented in Java, PHP and JavaScript use very similar identifiers "needle" and "haystack".

Quantitatively evaluating this hypothesis requires multiple implementations of the same problem in different programming languages, where we could compare identifier names. Luckily, RosettaCode provides just such a dataset. RosettaCode currently consists of 1,110 tasks, 305 draft tasks and includes 838 languages[3]. We collect the mined data[4] and study the same six languages (*i.e.*, Ruby, JavaScript, Java, Go, PHP, and Python) in the CodeXGLUE dataset. We get 15 cross-language pairs from six languages and measure identifier similarity between pairs of programs which solve the same problem in each language (*e.g.,* programs for graph diameter problem in Java and Ruby). For baselining, we also compare with a random pair (solving different problems) for the same two languages (*e.g.* graph diameter in Java, and SHA-hashing in Ruby). Fortunately, we found sufficient sample sizes for all our language pairs in RosettaCode. For example, for Java & Python we find 544 matched program pairs solving the same problem in both languages. We then take the 544 Java programs and randomly pair them with 544 other Python programs. Therefore, we have two groups of programs (*i.e.*, same program implemented in different languages and different programs implemented in different languages), and we check the similarity level between the two groups. Note that *size-unrestricted* random pairing may yield misleading results. Suppose we have a Java & Python program *matched* pair with 100 Java subtokens and 40 Python subtokens. Now, if we replace the matched python program with a random, bigger program (*e.g.*, 500 subtokens), we may have more chance of finding matched identifiers. Therefore, while choosing the random program, we try to ensure it has a similar length to the program it is replacing in the pair. We randomly select a program having the subtoken counts within a 5% length range (*e.g.*, 38-42 subtokens for a 40 subtoken program) of the removed one. Fortunately, in *99.25%* cases, we get at least one example within the 5% range. On the remaining instances, we select the program with the nearest subtoken count.

We measure identifier similarity thus:

(1) Remove all keywords, operators, and separators from the programs.
(2) Break all CamelCase and snake_case identifiers and keep only one copy of each sub token.
(3) Discard too-small programs with less than 5 sub-tokens.
(4) Calculate the mean Szymkiewicz-Simpson coefficient (overlap coefficient) [65] for both groups (*i.e.*, same program pair and random pair) of programs.
(5) Repeat this process across all 15 language pairs, for all program pairs.

Table 3 shows the common program pairs have 89%-235% *additional* **identifier** *overlap* compared to random program pairs. We compare the matched and random pair overlaps using the non-parametric Wilcoxon signed-rank test (AH: random has less overlap than matched). We observe that the null hypothesis is rejected, and Szymkiewicz-Simpson Overlap coefficient[5] is significantly higher for the common program pairs in all the cases. That indicates programs solving the same problem (even in different languages) are much more likely to use the same or similar identifier names.

Language pair	#of common programs	Overlap coefficient for random programs	Overlap coefficient for common programs	Overlap coefficient increased in %	Effect Size	p-value (adjusted)
Java & Python	544	0.10	**0.32**	+210.67%	0.747	<0.001
Java & Ruby	532	0.11	**0.31**	+174.97%	0.751	<0.001
Java & Javascript	411	0.13	**0.36**	+188.17%	0.774	<0.001
Java & Go	602	0.19	**0.36**	+89.24%	0.641	<0.001
Java & PHP	282	0.08	**0.28**	+235.01%	0.740	<0.001
Python & Ruby	538	0.11	**0.35**	+228.89%	0.780	<0.001
Python & Javascript	377	0.12	**0.34**	+190.09%	0.728	<0.001
Python & Go	601	0.13	**0.31**	+133.06%	0.664	<0.001
Python & PHP	267	0.09	**0.29**	+214.32%	0.679	<0.001
Ruby & Javascript	370	0.13	**0.35**	+167.02%	0.751	<0.001
Ruby & Go	571	0.12	**0.28**	+133.47%	0.724	<0.001
Ruby & PHP	262	0.09	**0.28**	+205.32%	0.716	<0.001
Javascript & Go	418	0.14	**0.29**	+110.96%	0.635	<0.001
Javascript & PHP	236	0.11	**0.29**	+175.03%	0.678	<0.001
Go & PHP	293	0.10	**0.23**	+121.25%	0.562	<0.001
Overall	6304	0.12	**0.31**	+158.94%	0.697	0

Table 3: ***Cross-language identifier similarity, when functionality is preserved***

We also calculate each pair's Jaccard index [31] (similarity coefficient) and find 112%-309% more similarity between common pairs than random ones, thus, giving essentially the same result. However, we prefer to report the detailed result using the overlap coefficient because Jaccard index can be affected by the differing verbosity of languages. For example, on average, Java, Python, and Ruby programs in RosettaCode have 29.45, 17.93, and 17.63 identifier subtokens. Java has higher subtokens compared to Python and Ruby because of the import statements, package naming etc. Therefore, Jaccard index between Java and Python will be lower than that of Python and Ruby even if the programs use very similar identifiers.

Finding 2. For a given problem, developers are likely to choose similar identifiers, even if coding in different languages.

In this section, we have presented evidence suggesting that a) identifiers are important for code summarization, that b) cross-language training is promising, and also that c) identifiers tend to be used in similar ways across languages. Taken together, these findings present a strong argument to try multilingual fine-tuning for SE tasks. Note that it is already well established that multilingual pre-training is helpful, and most BERT-style SE pre-trained models are multilingual [1, 18, 54, 55]. However, pre-training data are unsupervised and easy to collect. Preparing clean data for the supervised fine-tuning phase requires more time and attention. In this paper, our aim is to prove that multilingual training is not only effective in pre-training stage but also in fine-tuning stage for SE models, which is already found to be beneficial for *natural language* models [63].

3 BENCHMARK DATASETS AND TASKS

We evaluate the benefits of multilingual training in the context of several tasks, and associated datasets. In this section, we discuss the models and tasks used for our experiments.

[3]Last Accessed August, 2021
[4]https://github.com/acmeism/RosettaCodeData
[5]This is a measure of similarity like the Jaccard index; we use it here since sometimes the sizes of the programs are quite different. It's calculated as $\frac{|X \cap Y|}{min(|X|,|Y|)}$.

```java
public static int indexOf(ByteBuf needle, ByteBuf haystack) {
  // TODO: maybe use Boyer Moore for efficiency.
  int attempts = haystack.readableBytes() - needle.readableBytes() + 1;
  for (int i = 0; i < attempts; i++) {
    if (equals(needle, needle.readerIndex(),
      haystack, haystack.readerIndex() + i,
      needle.readableBytes())) {
      return haystack.readerIndex() + i;
    }
  }
  return -1;
}
```

(a) *Java*

```php
public static function indexOf(string $haystack, string $needle,
        int $offset=0):int
{
  $pos=self::strpos($haystack, $needle, $offset);
  return is_int($pos)?$pos:-1;
}
```

(b) *PHP*

```javascript
function indexOf(haystack, needle) {
  if (typeof haystack==='string')
    return haystack.indexOf(needle);
  for (let i=0, j=0, l=haystack.length, n=needle.length; i<l; i++) {
    if (haystack[i]===needle[j]) {
      j++;
      if (j===n) return i-j+1;
    }
    else {
      j=0;
    }
  }
  return -1;
}
```

(c) *JavaScript*

Figure 2: ***Usage of similar identifiers (e.g., needle, haystack) in "indexOf" function in different programming languages***

3.1 The Models

For our study of multilingual training, we adopt the BERT, or "foundation model" paradigm. Foundation models [13, 15, 17, 45, 56] have two stages: i) unsupervised pre-training with corpora at vast scale and ii) fine-tuning with a smaller volume of supervised data for the actual task. Foundation models currently hold state-of-the-art performance for a great many NLP tasks. BERT [17] style models have also been adapted for code, pre-trained on a huge, multilingual, corpora, and made available: CodeBERT and GraphCodeBERT are both freely available: both source code and pre-trained model parameters. While these models for code have thus far generally been fine-tuned monolingually, they provide an excellent platform for training experiments like ours, to measure the gains of multilingual fine-tuning. CodeBERT & GraphCodeBERT use a multi-layer bidirectional Transformer-based [64] architecture, and it is exactly as same as the RoBERTa [45], with 125M parameters; we explain them further below.

Pre-training The CodeBERT [18] dataset, has two parts: a matched-pairs part with 2.1M pairs of function and associated comment (NL-PL pairs) and 6.4M samples with just code. The code includes several programming languages. It was created by Hussain *et al.* [29]. CodeBERT model is pre-trained with two objectives (*i.e.*, Masked Language Modeling and Replaced Token Detection) on both parts. Mask language Modeling (MLM) is a widely applied and effective [17, 45] training objective where a certain number of (15%) tokens are masked out, and the model is asked to find those tokens. For CodeBERT training, Feng *et al.* apply this first objective only to bimodal data [18]. The second objective, Replaced Token Detection (RTD) [14], is a binary classification problem that is applied to both unimodal and bimodal data. Two data generators (*i.e.*, NL and PL) generate plausible alternatives for a set of randomly masked positions, and a discriminator is trained to determine whether a word is the original one or not. We note that CodeBERT pre-training is all about representation-learning: by learning to perform the task well, the model learns a good way to *encode* the text, which is helpful during the next, fine-tuning stage. The pre-training took about 12 hours on a machine with 16 NVIDIA V100 cards, and would have taken us very much longer, so we were grateful to be able to just download the estimated parameters.

Pre-training GraphCodeBERT GraphCodeBERT augments source-code with data flow, during pre-training. It uses a simple data flow graph (DFG) encoding a *where-the-value-comes-from* relation between variables [22]. The DFG nodes are variable occurrences, edges are value flow. GraphCodeBERT pretraining learns a joint representation of 1) the DFG structure, 2) DFG alignment with source code, and 3) the source code token sequences. GraphCodeBERT is therefore pre-trained with three training objectives (*i.e.*, Edge Prediction, Node Alignment, and MLM) on 2.3M functions (PL-NL pairs) from CodeSearchNet [29] dataset. For details see [22].

The pre-training+fine-tuning approach relies on VERY high capacity models, and are pre-trained over a large, multilingual corpus. Thus, even before fine-tuning, the models already know a lot about each language. Thus, fine-tuning on many languages should not negatively impact what the model knows about any one language. Thus we find that multilingual fine-tuning improves on monolingual fine-tuning in most cases. We believe our proposed approach would still consider the context surrounding the individual programming language even after multilingual training because these models have sufficient capacity to do so.

We now describe our tasks: in each, we describe the task, the dataset, and the multilingual fine-tuning approach (if applicable).

3.2 Code Summarization

The Task: as described earlier, the goal is to generate a NL summary given code in some PL.

The Dataset: There are several different code summarization datasets; we chose CodeXGLUE[6] [47], for two main reasons:

(1) CodeXGLUE is carefully de-duplicated [60]. Prior datasets like TL-CodeSum [28] have duplicates [60] in training, testing, and validation partitions. Duplication can inflate measured performance [5, 60].
(2) We need a multilingual dataset to prove the effectiveness of multilingual fine-tuning. None of the existing datasets [28, 40] is multilingual.

Table 4 presents the number of training, testing and validation instances for each language. in CodeXGLUE.

Model & Fine-tuning Feng *et al.* use a transformer-based encoder-decoder architecture for the code summarization task [18]. The

[6] CodeSearchNet [29] dataset is a standard benchmark, which has been incorporated into CodeXGLUE

Programming language	Training	Dev	Test	Candidate codes*
Ruby	24,927	1,400	1,261	4,360
JavaScript	58,025	3,885	3,291	13,981
Java	164,923	5,183	10,955	40,347
Go	167,288	7,325	8,122	28,120
PHP	241,241	12,982	14,014	52,660
Python	251,820	13,914	14,918	43,827

*Candidate codes are only used for code retrieval task

Table 4: CodeXGLUE ***dataset***

encoder is all ready well-trained in the pre-training stage; for fine-tuning, the encoder is primed with weights from pre-training. Now, the transformer model is given the input code token sequence and asked to generate the comment, as in the Neural Machine Translation (NMT) problem. We fine-tune using the CodeXGLUE paired samples. During fine-tuning, the decoder is trained auto-regressively, using next-token cross-entropy loss. Feng *et al.* use smooth BLEU-4 [44] for the evaluations of the models. Subsequently, We replace the pre-trained CodeBERT with pre-trained GraphCodeBERT in the encoder while evaluating the effectiveness of multilingual fine-tuning with GraphCodeBERT.

Why baseline with CodeBERT *for code summarization?* Feng *et al.* compare CodeBERT with other popular encoder-decoder based (*e.g.*, LSTM [61], Transformer [64], RoBERTa [45]) models; CodeBERT handily beats all of them [18]. Thus, CodeBERT is a good baseline to measure the value of multilingual finetuning. CodeBERT also does very well on prior datasets: using smoothed Sentence BLEU-4, we found that CodeBERT reaches 44.89 on TL-Codesum [28], and 32.92 on Funcom [40][7]. TL-Codesum has high degree of duplicates; we found that Funcom also does, but just in the comments. CodeXGLUE has very little duplication, which makes it more challenging, and also more reliable. Note that GraphCodeBERT *does not report* any performance on the code summarization task, and so we had to measure it.

3.3 Code Search

The Task Given a natural language query, find the semantically closest code sample from a large set of candidates. Vector-based information retrieval methods can be used here along with BERT-style encoders. CodeBERT was shown to perform quite well; the best published performance is reported by GraphCodeBERT [22] (CodeBERT augmented with graph representations). We study the value of multilingual fine-tuning for both CodeBERT and GraphCodeBERT (pre-training of both models was discussed earlier in Section 3.1).

The Dataset: Guo *et al.* adapt the same CodeSearchNet [29] dataset, with some additional data for candidate codes [22]. Note that it is basically the same dataset we used for code summarization except the candidate codes.

Model & Fine-tuning We use Guo *et al.*'s GraphCodeBERT model, which at the time of submission is the best performing model with code and parameters available, and so is fine-tunable. The fine-tuning data is code (PL) matched with (NL) comments, from CodeXGLUE. The pre-trained GraphCodeBERT embedding vector is calculated for each PL and NL part. During fine-tuning, Guo *et al.* take a minibatch of (say n) NL query vector, along with n (correct answers) PL answer vectors. n^2 dot products are calculated; the embedding vectors are then full-stack trained to give "1" normalized dot product for the matches, and "0" for the mis-matches. For the actual retrieval, GraphCodeBERT calculates the vector embedding of a given query, and simply retrieves candidates ranked by the dot-product distance from the query vector.

3.4 Method Name Prediction

The Task as introduced by Allamanis *et al.* [6] as the "extreme summarization" problem, the task is to predict the function name given the body.

The Dataset: We adapt the CodeXGLUE dataset by extracting the function name and asking the model to find the name given the function body. Following [6], the function names are broken into subtokens using BPE [59] (we've used BPE tokenization for all tasks). This problem then becomes very similar to code summarization.

Model & Fine-tuning Previously Code2Seq [7] and Code2Vec [8] have worked on this problem. All prior works [6–8] use a monolingual datasets, which are not suitable for our experiment. We use the same model we used for summarization, except we now learn to sequentially generate the method name, subtoken by subtoken. We use F1-score for the evaluation. For example, the function name "createLocal" is broken into two sub tokens (*i.e.*, create and Local), and the model predicts only "create". Hence, the precision, recall, and F1-score are 1.0, 0.5, and 0.66, respectively.

4 RESULTS

In this section, we evaluate multilingual fine-tuning for the baselines for the tasks enumerated above.

4.1 Code Summarization

We apply multilingual fine-tuning on the CodeXGLUE dataset. We first *replicate* the summarization task by (monolingually) fine-tuning the available pre-trained CodeBERT model for six languages[8]. We replicate the fine-tuning stage for 2 reasons:

(1) We want to account for any hardware or environmental bias (*e.g.*, we have a different set of GPUs than the original paper. We fine-tune with NVIDIA TITAN RTX, while Feng *et al.* [18] use NVIDIA Tesla V100).
(2) We use a pairwise two-sample statistical test (as described in [58], it is more precise than just comparing test-set summary statistics) to gauge differences. This requires a performance measurement for each test sample, which the repository did not include.

Our BLEU-4 numbers for monolingual training were close to reported numbers, with some differences; but we do obtain the same overall score (17.83) (table 5, leftmost 2 columns).

[7] As reported in [21, 60], measurement approaches vary across papers, and these numbers may differ from prior results: we use smoothed sentence BLEU-4 everywhere in our paper.

[8] We use the publicly available CodeBERT implementation and dataset, https://github.com/microsoft/CodeXGLUE/tree/main/Code-Text/code-to-text

We use the same, per-language test sets to compare monolingual and multilingual fine-tuning. The validation set, however, is a single multilingual one combining all the monolingual validation sets. Table 5 shows that multilingual fine-tuning improves performance, even for high-resource languages (with more than 100K training instances). With CodeBERT, multilingual fine-tuning gains 2.5%-17.5% over monolingual fine-tuning, for all languages, yielding a 6.90% overall improvement (4.48% weighted improvement)[9]. With the more advanced GraphCodeBERT, we see smaller gains, although the relative gains span a wide range.

We use a one-sided (AH: monolingual < multilingual) pairwise Wilcoxon signed-rank test (thus avoiding the corpus-level measurement pitfalls noted in [58]). Null hypothesis is rejected for all six languages, for CodeBERT. For GraphCodeBERT, it's rejected overall, and for every language; except for Javascript, where the p-value is 0.014 (all after B-H correction).

Thus our measurement indicates that multilingual fine-tuning provides a statistically significant improvement over monolingual training. We find rather low effect sizes using Cliff's Delta [48]. While we report the effect size for the sake of completeness, this is not a major concern: we note that *all gains are statistically highly significant*. We also emphasize that even the minor improvements provided here by multilingual training (which is broadly compatible with a range of settings) constitute a relevant and potentially widely useful result. Roy *et al* [58] have previously noted that small gains in BLEU-4 may not be perceptible to humans as increased text quality; nevertheless, we note that natural language translation (which is now widely used) attained high performance levels based on decades of incremental progress; this result and others below provide evidence that multilingual training could be an important step in the progress towards more useful automated tools. Finally, we note that BLEU-4 gains are higher for low-resource language (*e.g.*, 17.7% for Ruby), and lower for high-resource languages (*e.g.*, 2.5% for Python), as expected.

Comparing multi-lingual CodeBERT *with other models* Code summarization is widely studied—there are many models for this task; our specific focus here is to understand if multilingual fine-tuning provides benefits, using a high-quality token-sequence model and dataset. So we focus comparisons on the papers which report performance on CodeXGLUE dataset, and use a token-sequence inductive bias: comparing against all models is beyond the scope of this paper. We compare multi-lingual CodeBERT (*Polyglot*CodeBERT) and GraphCodeBERT (*Polyglot*GraphCodeBERT) with other models that have been published in peer-reviewed venues; among them, four apply pre-training strategies [1, 18, 45, 55]. We achieve the best overall performance (table 6), outperforming all the models, and for four specific languages (*i.e.*, Ruby, Java, Go and PHP).

There is one other system, CoTexT [54] which claims (in an unpublished, non-peer-reviewed report) better performance than us for just Python [54], but is worse overall. We will include it for comparison once it is published in a peer-reviewed venue.

This table also provides evidence supporting the effectiveness of multilingual fine-tuning.

[9]The CodeBERT paper simply averages the BLEU across languages to report the "overall" number; our *weighted average* weights each BLEU by the number of samples in that language.

4.2 Code Search

We study the gains from multilingual fine-tuning using two pre-trained models (*i.e.*,CodeBERT & GraphCodeBERT). We multilingually fine-tune both models using the publicly available code & dataset [10]. As we did for code summarization, we re-trained the baseline models, to get performance numbers for each case in the test set (to enable pairwise two-sample testing). We use the same test sets for both monolingual and multilingual training to evaluate our approach. During the training, GraphCodeBERT uses a matrix of dimension $|query| * |candidate\_codes|$. We could not use the full merged validation set (as we did for the code summarization task) because that makes the query and candidate code sets too large; the resulting matrix could not fit on our GPU server. We used a down-sampled validation set comprising six monolingual validation sets with 10K query and 50K candidate codes each. However, we did not face any issue while testing because we did not merge the test sets.

We report both the published values, and our replication; we need the replication to measure pairwise gains. Though CodeBERT and GraphCodeBERT both work on sequence of code tokens, GraphCodeBERT creates a rudimentary data-flow graph, once it's told the programming language.

Table 7 shows that multilingual fine-tuning improves the mean reciprocal rank for all languages except Go with CodeBERT. The improvement for Ruby, JavaScript, and Java are statistically significant. We found similar results for GraphCodeBERT exhibiting improvement for Ruby, JavaScript, Java, and Python; but with GraphCodeBERT both Go and PHP showed performance declines. However, overall, both showed statistically signficant improvements ($p < 0.001$); but the improvement for GraphCodeBERT (1.54%) is lower than CodeBERT (2.74%). Finally, we note that our numbers for CodeBERT differ from the performance reported for on the CodeXGLUE leaderboard. This is because CodeXGLUE benchmark uses only Python, and is based on a restricted setting where identifier names are left out. CodeXGLUE team argues that this abstraction enables them to stress-test the generalization ability of a model; however, here we consider an unmodified setting where someone gives an natural language query and wishes to find "natural" code with variable names intact.

4.3 Method Name Prediction

As for the previous two tasks, we try multilingual fine-tuning for method name prediction for CodeBERT. Here, too, we find evidence supporting the conclusion that multilingual training provides improvement for all the languages (Table 8). Non-parametric pairwise improvements are significant for Ruby, JavaScript, and Java. We also note observe relatively greater effect size for Ruby and JavaScript. Note that we achieve highest improvement for JavaScript because many functions therein are anonymous lambdas, since these functions have no names, they are not useful, and this diminishes available the JavaScript training set relative to other tasks (lambdas still have summaries, and can be used for other tasks). Therefore, multilingual fine-tuning increases the dataset diversity and boosts JavaScript method name prediction performance.

[10]https://github.com/microsoft/CodeBERT/tree/master/GraphCodeBERT/codesearch

Language	CodeBERT (reported)	CodeBERT (re-trained)	*Polyglot*CodeBERT	Improvement	Effect Size	p-value (adjusted)	GraphCodeBERT	*Polyglot*GraphCodeBERT	Improvement	Effect Size	p-value (adjusted)
Ruby	12.16	12.53	14.75	+17.72%	0.055	<0.001	12.62	**14.95**	+18.46%	0.055	<0.001
JS	14.90	13.86	**15.80**	+14.00%	0.016	<0.001	14.79	15.79	+6.76%	0.016	0.014
Java	17.65	18.72	**20.11**	+7.43%	0.016	<0.001	19.22	19.91	+3.59%	0.016	<0.001
Go	18.07	18.15	18.77	+3.42%	0.010	<0.001	18.40	**18.92**	+2.83%	0.010	<0.001
PHP	25.16	25.48	**26.23**	+2.94%	0.012	<0.001	25.45	26.15	+2.75%	0.012	<0.001
Python	**19.06**	18.25	18.71	+2.52%	0.022	<0.001	18.02	18.90	+4.88%	0.022	<0.001
Overall	17.83	17.83	19.06	+6.90%	0.016	<0.001	18.08	**19.10**	+5.64%	0.016	<0.001
Overall (weighted)	Not Reported	19.85	20.74	+4.48%			19.98	**20.76**	+3.90%		

*Evaluation criteria followed by CodeXGLUE [47] and CodeBERT [18]

Table 5: ***Effectiveness of multi-lingual fine-tuning for code summarization task. Note that p-values are B-H corrected***

4.4 Two Illustrative Examples

We used the same dataset for all tasks; for illustration, we show (Table 9) two test instances where all the tasks show improved performance from multilingual fine-tuning. In code summarization task, the monolingual fine-tuning scores 25 BLEU-4 in Example 1. CodeBERT produces a semantically wrong comment where multilingual fine-tuning generates the semantically correct solution. Note that the BLEU-4 is 84 for the second example because of the missing period in the gold standard (BLEU-4 is case-insensitive). Multilingual fine-tuning also helps the code search problem by increasing the MRR from 0.33 (Rank:3) to 1.00 (Rank:1). We also observe performance improvement from the method name prediction task. The gold standard consists of two sub tokens (*i.e.*, set and Values), and mono-lingual fine-tuning generates three (*i.e.*, set, Array, and Value), one of them is exact match. On the other hand, multilingual fine-tuning removes the extra "Array" subtoken and produces two subtokens(*i.e.*, set and Value) resulting in the F-score 0.50. We observe a similar result in example 2. Note that like BLEU-4, our method name prediction metric is also case-insensitive.

Finding 3. Multilingual fine-tuning is likely to increase diversity and help the models perform better than those trained with smaller mono-lingual datasets, especially for low-resource languages, irrespective of the task.

5 INTERPRETING RESULTS, AND THREATS

In this section we consider several issues that are relevant to the observed performance of multilingual training, such as model choice, dataset duplication, performance metrics, generalization, and different training strategies for the models.

Models	Overall	Ruby	JavaScript	Go	Python	Java	PHP
*Polyglot*GraphCodeBERT	**19.10**	**14.95**	15.79	**18.92**	18.90	19.91	26.15
*Polyglot*CodeBERT	19.06	14.75	15.80	18.77	18.71	**20.11**	**26.23**
ProphetNet-X [55]	18.54	14.37	**16.60**	18.43	17.87	19.39	24.57
PLBART [1]	18.32	14.11	15.56	18.91	**19.30**	18.45	23.58
GraphCodeBERT [22]	18.08	12.62	14.79	18.40	18.02	19.22	25.45
CodeBERT [18]	17.83	12.16	14.90	18.07	19.06	17.65	25.16
RoBERTa [45]	16.57	11.17	11.90	17.72	18.14	16.47	24.02
Transformer [64]	15.56	11.18	11.59	16.38	15.81	16.26	22.12
Seq2Seq [61]	14.32	9.64	10.21	13.98	15.93	15.09	21.08

Table 6: ***Comparison to existing models, on*** CodeXGLUE ***dataset***

5.1 Does multilingual fine-tuning help with other models?

There are several models, including CoTexT [54], ProphetNet-X [55], and PLBART [1] which report higher performance than CodeBERT [18] model for the code summarization task. The models for all these tasks were fine-tuned using monolingual datasets, so we might expect that multilingual fine-tuning should improve performance. These experiments would require a substantial investment of compute energy and is left for future work. We focused on CodeBERT (and also GraphCodeBERT on some tasks). We did some preliminary experiments with multilingual fine-tuning on PLBART. In our preliminary study, did see the same gains for low-resource language (Ruby, 5% gain). However, we found a 0.55% overall loss, which is inconsistent with what we observe with *Polyglot*CodeBERT (6.90% overall improvement) & *Polyglot*GraphCodeBERT (5.64% overall improvement). More study is needed.

Finding 4. Multilingual fine-tuning could benefit a broad range of models. We find gains for CodeBERT and GraphCodeBERT, but more data is required for other models.

5.2 Threats: Risk of data duplication?

Data duplication can lead to poor-quality estimates of performance, especially when data is duplicated across training & test; even duplication just within test data risks higher variance in the estimates. Allamanis finds that performance metrics are highly inflated when test data has duplicates, and advocates de-duplicating datasets, for more robust results [5]. Shi *et al.* also discusses the impact of duplication in code summarization task [60].

Sadly, there is a large amount of copied code on GitHUB [46]; inattentively combining different datasets harvested from GitHUB can lead to undesirable levels of duplication in the merged dataset. Fortnuately, CodeXGLUE is atually a *carefully de-duplicated* dataset; performance estimates therein are thus more robust. Combining *multilingual* data is unlikely to introduce the same kind of exact duplication in the dataset, because of syntax differences; There is a possibility of cross-language clones [53]; the study of this is left for future work.

Finding 5. Combining multilingual datasets is unlikely to cause exact duplication, because of syntax differences. More study is needed to study the effect of cross-language clones.

Language	CodeBERT (published) [22]	CodeBERT (re-trained)	*Polyglot*CodeBERT	Improvement	Effect Size	p-value (adjusted)	GraphCodeBERT (published) [22]	GraphCodeBERT (re-trained)	*Polyglot*GraphCodeBERT	Improvement	Effect Size	p-value (adjusted)
Ruby	0.679	0.677	0.732	+8.12%	0.072	<0.001	0.703	0.708	**0.738**	+4.24%	0.039	<0.001
JavaScript	0.620	0.616	0.643	+4.38%	0.034	<0.001	0.644	0.644	**0.660**	+2.48%	0.019	0.004
Java	0.676	0.676	0.697	+3.11%	0.026	<0.001	0.691	0.693	**0.710**	+2.45%	0.022	<0.001
Go	0.882	0.885	0.885	0%	-0.003	0.550	**0.897**	0.894	0.894	0%	-0.002	0.724
PHP	0.628	0.629	0.635	+0.95%	0.009	0.003	**0.649**	0.648	0.646	-0.31%	-0.002	0.904
Python	0.672	0.676	0.678	+0.30%	0.004	0.050	0.692	0.692	**0.695**	+0.43%	0.005	0.300
Overall*	0.693	0.693	0.712	+2.74%	0.013	<0.001	0.713	0.713	**0.724**	+1.54%	0.007	<0.001
Overall (weighted)	Not Reported	0.692	0.702	+1.42%			Not Reported	0.709	**0.715**	+0.80%		

*Evaluation criteria followed by GraphCodeBERT [22]

Table 7: *Effectiveness of multi-lingual fine-tuning for code search task. Note that p-values are BH-corrected*

Language	CodeBERT			*Polyglot*CodeBERT			F-Score Improvement	Effect Size	p-value (adjusted)
	Precision	Recall	F-Score	Precision	Recall	F-Score			
Ruby	0.44	0.40	0.41	0.53	0.49	**0.49**	20.59%	0.112	<0.001
JavaScript	0.30	0.24	0.26	0.45	0.40	**0.41**	59.00%	0.215	<0.001
Java	0.54	0.51	0.51	0.56	0.52	**0.52**	2.22%	0.016	<0.001
Go	0.54	0.52	0.52	0.56	0.53	0.52	1.67%	0.015	0.004
PHP	0.56	0.53	0.52	0.57	0.53	**0.53**	1.30%	0.009	0.004
Python	0.49	0.45	0.45	0.50	0.45	**0.46**	1.60%	0.011	0.002
Overall	0. 48	0.44	0.44	0.53	0.49	**0.49**	10.09%	0.024	<0.001
Overall (weighted)	0. 52	0.48	0.48	0.54	0.50	**0.50**	3.37%		

Table 8: *Effectiveness of multi-lingual fine-tuning for method naming task. Note that p-values are adjusted using Benjamini-Hochberg*

5.3 Threats: Other metrics?

Following CodeXGLUE benchmark recommendation, we evaluate the code summarization task with smooth sentence BLEU-4 [44] throughout this paper. However, other recognized metrics are are available (*e.g.*, ROUGE-L [43], METEOR [10]). Prior works [21, 58, 60] provide a careful analysis of the metrics, baselines, evaluations for code summarization task. Table 10 shows ROUGE-L and METEOR data; we find that multilingual fine-tuning increases the overall performance by 4.89% and 5.61% in ROUGE-L and METEOR, respectively. As with BLEU-4, we find that multilingual fine-tuning shows similar performance gains with these metrics. We find 0.3%-14.1% improvement in ROUGE-L and 1.2%-22.5% gains in METEOR (except for PHP, were we see a 0.17% *decline*, not statistically significant). We also see that Python shows the smallest improvement, not as strongly statistically significant. These metrics also indicate strong gains from multilingual training for low-resource and narrow-domain languages (*i.e.*, Ruby and JavaScript).

Finding 6. We observe performance improvement in all code summarization metrics with multilingual fine-tuning.

5.4 Monolingual minibatches? or multilingual?

While training deep neural networks with stochastic gradient descent, gradients (multivariate derivatives of loss *w.r.t* learnable parameters) are estimated over *mini-batches*, rather than calculating loss gradients over the entire training set; these estimates are used to adjust the weights in the network. Better choices of mini-batches could improve convergence behavior. With multilingual training, a natural question arises: is it better to sequentially interperse *monolingual* mini-batches (*e.g.,* first a Java minibatch, then Ruby minibatch and so on, before going back to Java?) or should we make each minibatch *per se* multilingual?

Example:1

```
//set the values from an Array
public void setValues* ( Array arr ) {
    //we omit intermediate lines to fit in the paper
    //original code here
}
```

Code Summarization

Models & comments	BLEU-4
Gold: set the values from an Array	NA
CodeBERT: Sets the values of the array .	25
*Polyglot*CodeBERT: Set the values from an array .	84

Code Search

Models	MRR
GraphCodeBERT	0.33
*Polyglot*GraphCodeBERT	1.00

Method Name Prediction

Models & method name	Sub tokens	F-Score
Gold: setValues	set Values	NA
CodeBERT: setArrayValue	set Array Value	0.40
*Polyglot*CodeBERT: setValue	set Value	0.50

Example:2

```
//Registers set injection point .
public void registerPetiteSetInjectionPoint* ( final String beanName, final String property ) {
    //we omit intermediate lines to fit in the paper
    //original code here
}
```

Code Summarization

Models & comments	BLEU-4
Gold: Registers set injection point .	NA
CodeBERT: Register a set of set InjectionPoint .	19
*Polyglot*CodeBERT: Register a set injection point .	60

Code Search

Models	MRR
GraphCodeBERT	0.50
*Polyglot*GraphCodeBERT	1.00

Method Name Prediction

Models & method name	Sub tokens	F-Score
Gold: registerPetiteSetInjectionPoint	register Pet ite Set In jection Point	NA
CodeBERT: addPropertyInjectionPoint	add Property In jection Point	0.50
*Polyglot*CodeBERT: setPropertyInjectionPoint	set Property In jection Point	0.57

*** "registerPetiteSetInjectionPoint" & "setValues" tokens are abstracted for method name prediction task*

Table 9: *Examples exhibiting the effectiveness of multilingual training*

In the previous experiments, we had randomly sort the dataset; hence, our mini-batches are also multilingual. So we deliberately tried sequentially monolingual minibatching during multilingual fine-tuning. We find that sequentially monolingual minibatch training appears to somewhat degrade performance: we observe the overall performance goes down by 1.05%. However, the change is not statistically significant for any language. We omit the actual numerical details, for space reasons, since we didn't find any strong results in either direction.

Language	ROUGE-L				METEOR			
	*Polyglot*CodeBERT	Improve.*	Effect Size	p-value (adjusted)	*Polyglot*CodeBERT	Improve.*	Effect Size	p-value (adjusted)
Ruby	24.36	+14.10%	0.087	<0.001	21.96	+22.54%	0.125	<0.001
JavaScript	24.30	+7.05%	0.022	<0.001	21.59	+11.40%	0.030	<0.001
Java	34.89	+3.32%	0.020	<0.001	31.73	+4.41%	0.020	<0.001
Go	37.36	+2.69%	0.024	<0.001	30.28	+3.73%	0.023	<0.001
PHP	38.81	+0.34%	-8.65E-05	0.508	35.52	-0.17%	-0.003	0.779
Python	32.86	+1.86%	0.015	<0.001	27.75	+1.24%	0.004	0.033
Overall	32.10	+4.89%	0.016	<0.001	28.14	+5.61%	0.013	<0.001
Overall (weighted)	34.82	+2.24%			30.52	+2.59%		

*Improvement reported over CodeBERT

Table 10: ***Performance improvement in ROUGE-L and METEOR for code summarization task***

Finding 7. We don't find any clear difference between multilingual mini batches and (interspersed) monolingual mini batches.

5.5 Multilingual model as pre-trained model

Our findings provide evidence supporting the claim that a multilingual fine-tuned model is effective for code summarization task, which outperforms all the models trained with monolingual datasets. Could this this improved multilingual model further benefit from a secondary, *monolingual* fine-tuning exercise, where it receives specialized fine-tuning for each language separately? To evaluate this intriguing and promising idea, we load the model with the weights from multilingual fine-tuning, and fine-tune it, again, for each individual language. Table 11 shows that We found some minor performance improvement for JavaScript and Python. However, the performance goes down for other languages. We do not find evidence that a secondary, monolingual fine-tuning is helpful; further work is needed to understand why, and perhaps develop other ways this idea might yield further improvement.

Language	*Polyglot*CodeBERT	*Polyglot*CodeBERT as pre-training	Improvement	Effect Size	p-value (adjusted)
Ruby	**14.75**	14.58	-1.15%	-0.016	0.303
JS	15.80	**16.47**	+4.24%	0.024	<0.001
Java	**20.11**	19.81	-1.49%	-0.003	0.303
Go	**18.77**	17.97	-4.26%	-0.012	<0.001
Php	**26.23**	25.52	-2.71%	-0.017	<0.001
Python	18.71	**18.83**	+0.64%	0.010	<0.001
Overall	**19.06**	18.86	-1.05%	-0.003	0.005
Overall (weighted)	**20.74**	20.43	-1.47%		

Table 11: ***Multilingual model as pre-trained model***

Finding 8. We don't find evidence that applying a secondary, mono-lingual fine-tuning provides benefits for all languages.

6 RELATED WORK

Code summarization: Code summarization has recently been a hot topic. More than 30 papers have been published in the last five years that follow some form of encoder-decoder architecture [58]. Several works [21, 58, 60] discuss the evaluations, metrics, and baselining. Roy *et al.* show that metric improvements of less than 2 points do not guarantee systematic improvements in summarization and are not reliable as proxies of human evaluation [58]. We find more than 2 points of improvement for Ruby and almost 2 points of improvement for JavaScript. We observe less than 2 points in other languages. It should also be noted that we *don't* use the corpus-level metrics which Roy *et al.* show is problematic; we use pairwise comparisons on the test-sets. Finally, we note that progress in both code & NLP occurs in small steps over decades, and innovations (especially ones that could cumulate with others) such as ours can be an important part of research community's long-term pursuit of practically relevant performance improvements.

Pre-trained models [1, 18, 45, 54, 55] are proven to be more effective than prior models. Different pre-trained models are trained with the different pre-trained objectives even though fine-tuning steps are almost similar for all the models. As discussed earlier in Section 3.1, CodeBERT is an encoder model, pre-trained with MLM and Repace Token Detection objectives. Unlike CodeBERT, PLBART [1] is an encoder-decoder model which is trained as a denoising auto-encoder. Though all the models are pre-trained with different training objectives, there is one thing common among all the models: using Transformers as core architecture.

Parvez *et al.* very recently present an approach that augments training data using relevant code or summaries retrieved from a database (*e.g.*, GitHub, Stack Overflow) [52]. They apply this approach on monolingual Java and Python datasets from CodeXGLUE and claim gains over *Polyglot*CodeBERT & *Polyglot*GraphCodeBERT for code summarization. *Prima facie*, multilingual fine-tuning is complementary to their approach; this needs to be studied.

Code retrieval and method name prediction: Code retrieval is also getting attention recently. There are multiple datasets for this task. CodeXGLUE introduces a monolingual python dataset (taken initially from CodeSearchNet) abstracting the function names and variables. Guo *et al.* modify the multilingual CodeSearchNet dataset and achieve state-of-the-art performance on this task. However, using multilingual training, we show that both CodeBERT and GraphCodeBERT can be improved. There is one other very recent paper, CLSEBERT [68] which reports (in an unpublished, non-peer-reviewed report) better performance than us in all languages except Ruby. We could not show the effectiveness of multilingual training on CLSEBERT because the authors have not released the code implementation yet. Note that like code summarization, we focus only on the work using CodeSearchNet multilingual dataset.

CodeSearchNet dataset can be easily adapted to method name prediction task. Several earlier works address method name prediction, in a Java-only such as Code2Seq [7], Allamanis [6]. They all use a conventional single-stage machine-learning approach (no pre-training + fine-tuning). Our goal here is to simply demonstrate that multilingual fine-tuning improves upon monolingual fine-tuning for the method-naming task, so we demonstrate using CodeBERT. Our numbers are roughly comparable with previously reported results, but we cannot make a precise comparison because of differences in subtokenization, and because our datasets are multilingual whereas previous work has largely been monolingual. We are simply arguing here our data suggests that multilingual fine-tuning is broadly beneficial in different tasks.

It would certainly be interesting to use same-domain data for fine-tuning. For example, summarizing methods in Android apps

might work better if trained on Android app corpora; however curated, domain-specific datasets for each domain are needed, and may not always be possible, depending on the domain. However, cross-language data is already available, and we show that it does indeed help improve performance! The effect of domain-specific corpora is left for future work.

7 CONCLUSION

We began this paper with three synergistic observations: *First,* when solving the *same problem*, even in different programming languages, programmers are more likely to use similar identifiers (than when solving *different* problems). *Second,* identifiers appear to be relatively much more important than syntax markers when training machine-learning models to perform code summarization. *Third,* we find that quite often a model trained in one programming language achieves surprisingly good performance on a test set in a different language, sometimes even surpassing a model trained on the same language as the test set! Taken together, these findings suggest that pooling data across languages, thus creating *multilingual* training sets, could improve performance for any language, particularly perhaps languages with limited resources, as has been found in Natural-language processing [16, 23, 57, 63]. We test this theory, using two BERT-style models, CodeBERT, and GraphCodeBERT, with encouraging results.

Foundation models [12] are currently achieving best-in-class performance for a wide range of tasks in both natural language and code. The models work in 2 stages, first "pre-training" to learn statistics of language (or code) construction from very large-scale corpora in a self-supervised fashion, and then using smaller labeled datasets to "fine-tune" for specific tasks. We adopt the multilingual CodeXGLUE dataset, and the pre-trained CodeBERT and GraphCodeBERT models, and study the value of multilingual fine-tuning for a variety of tasks. We find evidence suggesting that multilingual fine-tuning is broadly beneficial in many settings. Our findings suggest that multilingual training could provide added value in broad set of settings, and merits further study.

Acknowledgements: This material is based upon work supported by the U.S. National Science Foundation under Grant Nos. 1414172, and 2107592. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. Ahmed was also supported by UC Davis College of Engineering Dean's Distinguished Fellowship.

REFERENCES

[1] Wasi Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. 2021. Unified Pre-training for Program Understanding and Generation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Online, 2655–2668. https://www.aclweb.org/anthology/2021.naacl-main.211

[2] Wasi Uddin Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. 2020. A Transformer-based Approach for Source Code Summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*.

[3] Toufique Ahmed, Premkumar Devanbu, and Anand Ashok Sawant. 2021. Learning to Find Usage of Library Functions in Optimized Binaries. *IEEE Transactions on Software Engineering* (2021). https://doi.org/10.1109/TSE.2021.3106572

[4] Toufique Ahmed, Noah Rose Ledesma, and Premkumar Devanbu. 2021. SYNFIX: Automatically Fixing Syntax Errors using Compiler Diagnostics. arXiv:2104.14671 [cs.SE]

[5] Miltiadis Allamanis. 2019. The adverse effects of code duplication in machine learning models of code. In *Proceedings of the 2019 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*. 143–153.

[6] Miltiadis Allamanis, Hao Peng, and Charles Sutton. 2016. A convolutional attention network for extreme summarization of source code. In *International conference on machine learning*. PMLR, 2091–2100.

[7] Uri Alon, Shaked Brody, Omer Levy, and Eran Yahav. 2019. code2seq: Generating Sequences from Structured Representations of Code. In *International Conference on Learning Representations*. https://openreview.net/forum?id=H1gKYo09tX

[8] Uri Alon, Meital Zilberstein, Omer Levy, and Eran Yahav. 2019. code2vec: Learning distributed representations of code. *Proceedings of the ACM on Programming Languages* 3, POPL (2019), 1–29.

[9] Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George Foster, Colin Cherry, et al. 2019. Massively multilingual neural machine translation in the wild: Findings and challenges. *arXiv preprint arXiv:1907.05019* (2019).

[10] 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*. 65–72.

[11] Eeshita Biswas, Mehmet Efruz Karabulut, Lori Pollock, and K Vijay-Shanker. 2020. Achieving Reliable Sentiment Analysis in the Software Engineering Domain using BERT. In *2020 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 162–173.

[12] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the Opportunities and Risks of Foundation Models. *arXiv preprint arXiv:2108.07258* (2021).

[13] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165* (2020).

[14] Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555* (2020).

[15] Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. *Advances in Neural Information Processing Systems* 32 (2019), 7059–7069.

[16] Raj Dabre, Chenhui Chu, and Anoop Kunchukuttan. 2020. A survey of multilingual neural machine translation. *ACM Computing Surveys (CSUR)* 53, 5 (2020), 1–38.

[17] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).

[18] 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 *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*. 1536–1547.

[19] Shuzheng Gao, Cuiyun Gao, Yulan He, Jichuan Zeng, Lun Yiu Nie, and Xin Xia. 2021. Code Structure Guided Transformer for Source Code Summarization. *arXiv preprint arXiv:2104.09340* (2021).

[20] Ekaterina Garmash and Christof Monz. 2016. Ensemble learning for multi-source neural machine translation. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*. 1409–1418.

[21] David Gros, Hariharan Sezhiyan, Prem Devanbu, and Zhou Yu. 2020. Code to Comment ?Translation?: Data, Metrics, Baselining & Evaluation. In *2020 35th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 746–757.

[22] Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, LIU Shujie, Long Zhou, Nan Duan, Alexey Svyatkovskiy, Shengyu Fu, et al. 2020. GraphCodeBERT: Pre-training Code Representations with Data Flow. In *International Conference on Learning Representations*.

[23] Thanh-Le Ha, Jan Niehues, and Alexander Waibel. 2016. Toward multilingual neural machine translation with universal encoder and decoder. *arXiv preprint arXiv:1611.04798* (2016).

[24] Sakib Haque, Alexander LeClair, Lingfei Wu, and Collin McMillan. 2020. Improved automatic summarization of subroutines via attention to file context. In *Proceedings of the 17th International Conference on Mining Software Repositories*. 300–310.

[25] Masum Hasan, Tanveer Muttaqueen, Abdullah Al Ishtiaq, Kazi Sajeed Mehrab, Md Haque, Mahim Anjum, Tahmid Hasan, Wasi Uddin Ahmad, Anindya Iqbal, and Rifat Shahriyar. 2021. CoDesc: A Large Code-Description Parallel Dataset. *arXiv preprint arXiv:2105.14220* (2021).

[26] Xing Hu, Ge Li, Xin Xia, David Lo, and Zhi Jin. 2018. Deep code comment generation. In *2018 IEEE/ACM 26th International Conference on Program Comprehension (ICPC)*. IEEE, 200–20010.

[27] Xing Hu, Ge Li, Xin Xia, David Lo, and Zhi Jin. 2020. Deep code comment generation with hybrid lexical and syntactical information. *Empirical Software Engineering* 25, 3 (2020), 2179–2217.

[28] Xing Hu, Ge Li, Xin Xia, David Lo, Shuai Lu, and Zhi Jin. 2018. Summarizing source code with transferred api knowledge. (2018).

[29] 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).

[30] Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, and Luke Zettlemoyer. 2016. Summarizing source code using a neural attention model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 2073–2083.

[31] Paul Jaccard. 1901. Étude comparative de la distribution florale dans une portion des Alpes et des Jura. *Bull Soc Vaudoise Sci Nat* 37 (1901), 547–579.

[32] Kevin Jesse, Premkumar T Devanbu, and Toufique Ahmed. 2021. Learning type annotation: is big data enough?. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1483–1486.

[33] Nan Jiang, Thibaud Lutellier, and Lin Tan. 2021. CURE: Code-Aware Neural Machine Translation for Automatic Program Repair. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. IEEE, 1161–1173.

[34] Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. 2017. Google's multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association for Computational Linguistics* 5 (2017), 339–351.

[35] Aditya Kanade, Petros Maniatis, Gogul Balakrishnan, and Kensen Shi. 2020. Learning and evaluating contextual embedding of source code. In *International Conference on Machine Learning*. PMLR, 5110–5121.

[36] Surafel M Lakew, Mauro Cettolo, and Marcello Federico. 2018. A comparison of transformer and recurrent neural networks on multilingual neural machine translation. *arXiv preprint arXiv:1806.06957* (2018).

[37] Dawn Lawrie, Christopher Morrell, Henry Feild, and David Binkley. 2007. Effective identifier names for comprehension and memory. *Innovations in Systems and Software Engineering* 3, 4 (2007), 303–318.

[38] Alexander LeClair, Aakash Bansal, and Collin McMillan. 2021. Ensemble Models for Neural Source Code Summarization of Subroutines. *arXiv preprint arXiv:2107.11423* (2021).

[39] Alexander LeClair, Sakib Haque, Lingfei Wu, and Collin McMillan. 2020. Improved code summarization via a graph neural network. In *Proceedings of the 28th International Conference on Program Comprehension*. 184–195.

[40] Alexander LeClair, Siyuan Jiang, and Collin McMillan. 2019. A neural model for generating natural language summaries of program subroutines. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 795–806.

[41] Boao Li, Meng Yan, Xin Xia, Xing Hu, Ge Li, and David Lo. 2020. DeepCommenter: a deep code comment generation tool with hybrid lexical and syntactical information. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1571–1575.

[42] Zheng Li, Yonghao Wu, Bin Peng, Xiang Chen, Zeyu Sun, Yong Liu, and Deli Yu. 2021. SeCNN: A semantic CNN parser for code comment generation. *Journal of Systems and Software* 181 (2021), 111036.

[43] Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*. 74–81.

[44] Chin-Yew Lin and Franz Josef Och. 2004. Orange: a method for evaluating automatic evaluation metrics for machine translation. In *COLING 2004: Proceedings of the 20th International Conference on Computational Linguistics*. 501–507.

[45] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).

[46] Cristina V Lopes, Petr Maj, Pedro Martins, Vaibhav Saini, Di Yang, Jakub Zitny, Hitesh Sajnani, and Jan Vitek. 2017. DéjàVu: a map of code duplicates on GitHub. *Proceedings of the ACM on Programming Languages* 1, OOPSLA (2017), 1–28.

[47] 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 (2021).

[48] Guillermo Macbeth, Eugenia Razumiejczyk, and Rubén Daniel Ledesma. 2011. Cliff's Delta Calculator: A non-parametric effect size program for two groups of observations. *Universitas Psychologica* 10, 2 (2011), 545–555.

[49] Junayed Mahmud, Fahim Faisal, Raihan Islam Arnob, Antonios Anastasopoulos, and Kevin Moran. 2021. Code to Comment Translation: A Comparative Study on Model Effectiveness & Errors. *arXiv preprint arXiv:2106.08415* (2021).

[50] 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 (ICSE)*. IEEE, 336–347.

[51] Preslav Nakov and Hwee Tou Ng. 2009. Improved statistical machine translation for resource-poor languages using related resource-rich languages. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*. 1358–1367.

[52] Md Rizwan Parvez, Wasi Uddin Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. 2021. Retrieval Augmented Code Generation and Summarization. *arXiv preprint arXiv:2108.11601* (2021).

[53] Daniel Perez and Shigeru Chiba. 2019. Cross-language clone detection by learning over abstract syntax trees. In *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*. IEEE, 518–528.

[54] Long Phan, Hieu Tran, Daniel Le, Hieu Nguyen, James Anibal, Alec Peltekian, and Yanfang Ye. 2021. CoTexT: Multi-task Learning with Code-Text Transformer. *arXiv preprint arXiv:2105.08645* (2021).

[55] Weizhen Qi, Yeyun Gong, Yu Yan, Can Xu, Bolun Yao, Bartuer Zhou, Biao Cheng, Daxin Jiang, Jiusheng Chen, Ruofei Zhang, et al. 2021. ProphetNet-X: Large-Scale Pre-training Models for English, Chinese, Multi-lingual, Dialog, and Code Generation. *arXiv preprint arXiv:2104.08006* (2021).

[56] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683* (2019).

[57] Surangika Ranathunga, En-Shiun Annie Lee, Marjana Prifti Skenduli, Ravi Shekhar, Mehreen Alam, and Rishemjit Kaur. 2021. Neural Machine Translation for Low-Resource Languages: A Survey. *arXiv preprint arXiv:2106.15115* (2021).

[58] Devjeet Roy, Sarah Fakhoury, and Venera Arnaoudova. 2021. Reassessing automatic evaluation metrics for code summarization tasks. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1105–1116.

[59] Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909* (2015).

[60] Ensheng Shi, Yanlin Wang, Lun Du, Junjie Chen, Shi Han, Hongyu Zhang, Dongmei Zhang, and Hongbin Sun. 2021. Neural Code Summarization: How Far Are We? *arXiv preprint arXiv:2107.07112* (2021).

[61] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*. 3104–3112.

[62] Xu Tan, Yi Ren, Di He, Tao Qin, Zhou Zhao, and Tie-Yan Liu. 2019. Multilingual neural machine translation with knowledge distillation. *arXiv preprint arXiv:1902.10461* (2019).

[63] Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. Multilingual translation with extensible multilingual pretraining and finetuning. *arXiv preprint arXiv:2008.00401* (2020).

[64] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.

[65] MK Vijaymeena and K Kavitha. 2016. A survey on similarity measures in text mining. *Machine Learning and Applications: An International Journal* 3, 2 (2016), 19–28.

[66] 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*. 397–407.

[67] Wenhua Wang, Yuqun Zhang, Yulei Sui, Yao Wan, Zhou Zhao, Jian Wu, Philip Yu, and Guandong Xu. 2020. Reinforcement-learning-guided source code summarization via hierarchical attention. *IEEE Transactions on software Engineering* (2020).

[68] Xin Wang, Yasheng Wang, Pingyi Zhou, Meng Xiao, Yadao Wang, Li Li, Xiao Liu, Hao Wu, Jin Liu, and Xin Jiang. 2021. CLSEBERT: Contrastive Learning for Syntax Enhanced Code Pre-Trained Model. *arXiv preprint arXiv:2108.04556* (2021).

[69] Yanlin Wang, Ensheng Shi, Lun Du, Xiaodi Yang, Yuxuan Hu, Shi Han, Hongyu Zhang, and Dongmei Zhang. 2021. CoCoSum: Contextual Code Summarization with Multi-Relational Graph Neural Network. *arXiv preprint arXiv:2107.01933* (2021).

[70] Bolin Wei, Ge Li, Xin Xia, Zhiyi Fu, and Zhi Jin. 2019. Code generation as a dual task of code summarization. *arXiv preprint arXiv:1910.05923* (2019).

[71] Bolin Wei, Yongmin Li, Ge Li, Xin Xia, and Zhi Jin. 2020. Retrieve and refine: exemplar-based neural comment generation. In *2020 35th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 349–360.

[72] Guang Yang, Xiang Chen, Jinxin Cao, Shuyuan Xu, Zhanqi Cui, Chi Yu, and Ke Liu. 2021. ComFormer: Code Comment Generation via Transformer and Fusion Method-based Hybrid Code Representation. *arXiv preprint arXiv:2107.03644* (2021).

[73] Huang Yuchao, Wei Moshi, Wang Song, Wang Junjie, and Wang Qing. 2021. Yet Another Combination of IR-and Neural-based Comment Generation. *arXiv preprint arXiv:2107.12938* (2021).

[74] Jian Zhang, Xu Wang, Hongyu Zhang, Hailong Sun, and Xudong Liu. 2020. Retrieval-based neural source code summarization. In *2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE)*. IEEE, 1385–1397.

[75] Ting Zhang, Bowen Xu, Ferdian Thung, Stefanus Agus Haryono, David Lo, and Lingxiao Jiang. 2020. Sentiment Analysis for Software Engineering: How Far Can Pre-trained Transformer Models Go?. In *2020 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 70–80.

[76] Barret Zoph and Kevin Knight. 2016. Multi-source neural translation. *arXiv preprint arXiv:1601.00710* (2016).