

Optimizing Datasets for Code Summarization: Is Code-Comment Coherence Enough?

*Original*

Optimizing Datasets for Code Summarization: Is Code-Comment Coherence Enough? / Vitale, A.; Mastropaolo, A.; Oliveto, R.; Di Penta, M.; Scalabrino, S.. - (2025), pp. 237-249. ( 2025 IEEE/ACM 33rd International Conference on Program Comprehension (ICPC) Ottawa (CAN) 27-28 April 2025) [10.1109/ICPC66645.2025.00033].

*Availability:*

This version is available at: 11583/3002676 since: 2025-09-02T11:08:04Z

*Publisher:*

IEEE

*Published*

DOI:10.1109/ICPC66645.2025.00033

*Terms of use:*

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

*Publisher copyright*

IEEE postprint/Author's Accepted Manuscript

©2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

# Optimizing Datasets for Code Summarization: Is Code-Comment Coherence Enough?

Antonio Vitale<sup>\*‡</sup>, Antonio Mastropaolo<sup>†</sup>, Rocco Oliveto<sup>‡</sup>, Massimiliano Di Penta<sup>§</sup>, and Simone Scalabrino<sup>§</sup>

<sup>\*</sup>Politecnico di Torino, Italy, antonio.vitale@polito.it

<sup>†</sup>William & Mary, USA, amastropaolo@wm.edu

<sup>‡</sup>University of Molise, Italy, {rocco.oliveto, simone.scalabrino}@unimol.it

<sup>§</sup>University of Sannio, Italy, dipenta@unisannio.it

**Abstract**—Automated code summarization is a long-standing goal for code comprehension. This task automatically generates documentation using a given method. Deep Learning (DL)-based approaches have been proven beneficial for various software engineering (SE) tasks, including this one. Most state-of-the-art datasets for code summarization are automatically mined from GitHub and, thus, might contain erroneous or sub-optimal examples. Previous work showed that using a simple rule-based approach for removing noisy instances allows for a tangible reduction of the training set size while not reducing the effectiveness of the trained models. Motivated by this finding, we conjecture that it is possible to further reduce the dataset size by removing instances that contain different issues. In this paper, we explore the extent to which code-comment coherence, a specific quality attribute of code summaries, can be used to optimize code summarization datasets. Specifically, we hypothesize that removing incoherent code-comment pairs might positively impact the effectiveness of the models. To do this, we rely on SIDE, a recently introduced metric for code-summary coherence. We examine multiple selectivity levels of training instances from two state-of-the-art datasets (TL-CodeSum and Funcom) and evaluate the resulting models on three manually curated test sets. The results show that even halving the training set sizes does not significantly affect the model’s ability to generate summaries. However, when comparing the most restrictive selection strategy with a simpler one that randomly selects the training instances, we observe that the resulting accuracy of the model also does not change. This result suggests that (i) current datasets contain many irrelevant examples, and (ii) different quality attributes should be explored for optimizing code summarization datasets.

**Index Terms**—Code Summarization, Data Quality, Code-Comment Coherence, Empirical Study

## I. INTRODUCTION

Writing comments and keeping them up to date during software maintenance and evolution is a challenging task and requires effort [1], [2], [3], [4]—yet, high-quality comments and documentation are essential for understanding code [5], [6]. Thus, automatically generating code comments from code has always been a long-lasting dream of developers and practitioners willing to bolster program comprehension in a cost-effective way [7], [8], [9], [10], [11], [12]. A dream that came reality, thanks to the recent advancements in Deep Learning (DL) models and particularly Large Language Models (LLMs) that pushed the boundaries of Software Engineering (SE) automation to the next level.

```
public HashSet getCommandResultsRootFeatures() {
    HashSet rootFeatureSet = new HashSet();
    Feature belowSplitRoot = null;
    Feature aboveSplitRoot = null;
    if (belowSplitTranscript != null) {
        belowSplitRoot = belowSplitTranscript.getRootFeature();
        rootFeatureSet.add(belowSplitRoot);
    }
    if (aboveSplitTranscript != null) {
        aboveSplitRoot = aboveSplitTranscript.getRootFeature();
        if (aboveSplitRoot != belowSplitRoot)
            rootFeatureSet.add(aboveSplitRoot);
    }
    return rootFeatureSet;
}
```

Invoked AFTER the command is executed.

SIDE SCORE: -0.90

Fig. 1. Example of a  $\langle code, summary \rangle$  pair exhibiting a misalignment after the coarse-grained filtering by CAT.

Modern DL-based approaches based on Transformers [13], [14], [15] rely on transfer learning. First, a basic model is pre-trained to acquire knowledge of the programming language. Then, several specialized models can be fine-tuned from it to tackle specific tasks. This approach has been proven effective for several SE tasks, including code generation [16], [17], [18], [19], program repair [20], [21], [22], [23], and, indeed, code summarization [24], [25], [10], [12], [26], [11].

Fine-tuning a DL model still requires plenty of examples, which can hardly be produced or curated manually. Thus, most state-of-the-art datasets for training DL models for coding tasks (including code summarization) are built by automatically mining open-source software repositories.

However, datasets created by mining software repositories tend to be noisy, containing several low-quality instances [27], [28], [29]. More specifically, Shi *et al.* [30] recently proposed a heuristic-based approach named CAT that can automatically clean up low-quality  $\langle summary, code \rangle$  pairs from code summarization datasets. This approach removes or fixes structural quality issues at comment- (e.g., commented code) and code-level (e.g., empty methods). The results of their empirical study show that removing noisy instances does not reduce and even improves the model’s ability to produce meaningful code summaries. Even though CAT streamlines coarse-grained filtering, the  $\langle code, summary \rangle$  pairs that are not discarded can still have low quality and thus negatively impact the ability of the model to produce good code summaries. Let us consider

the instance in Fig. 1 from the state-of-the-art *Funcom* dataset. While no structural problem exists (CAT does not discard it), the instance is characterized by a summary completely incoherent with the source code. Exposing the model to instances with such inconsistencies can cause hallucinations [31], which may lead to decreased performance.

For this reason, we conjecture that a finer-grained selection of  $\langle \text{code}, \text{summary} \rangle$  pairs in code summarization datasets is not only possible, but even desirable to (i) reduce the training time, and (ii) possibly improve the model effectiveness. More specifically, we conjecture that code-comment coherence, which is a well-known quality attribute of  $\langle \text{code}, \text{summary} \rangle$  pairs [32], might play a crucial role in achieving this goal.

In this paper, we present an empirical investigation in which we study the impact of fine-grained filtering based on code-comment coherence on code summarization models in terms of *effectiveness* (i.e., correctness of the inferences) and *efficiency* (i.e., training time). To measure the code-comment coherence, we rely on SIDE, a metric recently introduced by Mastropaolo *et al.* [33]. The authors show that SIDE strongly correlates with human evaluations of summary quality, surpassing established metrics including BLEU [34] and ROUGE [35]. We consider two state-of-the-art datasets for code summarization, i.e., TL-CodeSum [24] and Funcom [9], already filtered with CAT [30]. Then, we further filter the instances in terms of SIDE value by considering different thresholds (i.e.,  $\{0.5, 0.6, 0.7, 0.8, 0.9\}$ ). Finally, we use each resulting training set (including the original one) to fine-tune *CodeT5+* [36].

We test the models on two manually-curated datasets from the literature [37], [38]. We observe that reducing the size of the training set, even with the most restrictive filter ( $\text{SIDE}_{0.9}$ , which selects  $\sim 50\%$  of the instances), has a negligible impact on the model’s effectiveness. On the other hand, reducing the number of training instances results in a significantly lower training time (up to  $\sim 111$  saved hours). To further validate our original hypothesis that code-comment coherence is a suitable quality attribute for filtering instances in code summarization datasets, we compared the most restrictive filter ( $\text{SIDE}_{0.9}$ ) with a filter that keeps the same number of instances, but by simply choosing them randomly. Surprisingly, we observed that the random filter achieves negligibly worse results than  $\text{SIDE}_{0.9}$ .

Our results provide two clear insights. First, code-comment coherence is marginally important for selecting suitable instances for code summarization. Second, regardless of this, removing instances (even randomly!) does not impact the effectiveness of code summarization models. This suggests that additional code-comment quality attributes should be investigated, and researchers should prioritize relevance over quantity to better select the most informative instances to build code summarization datasets.

The paper is organized as follows. Section II provides backgrounds on the existing selection technique for code summarization (CAT [30]) and overviews SIDE [33]. Section III details the study definition and planning. Results are reported in Section IV, while Section V discusses its implications, and

Section VI the threats to its validity. Section VII discusses related work about data quality. Finally, Section VIII concludes the paper and outlines directions for future work.

## II. SELECTION STRATEGIES FOR CODE SUMMARIZATION

In this section, we provide backgrounds about (i) a state-of-the-art strategy for repairing or removing poor instances from code summarization datasets (CAT), and (ii) the SIDE metric, which we use to streamline a data-centric, quality-aware instance filtering.

### A. Code-comment cleAning Tool

CAT (**C**ode-comment **cle**Aning **T**ool) is an approach and tool by Shi *et al.* [30] to detect and handle noisy instances given the pairs of  $\langle \text{code}, \text{summary} \rangle$  from code-summarization datasets. The development of CAT was preceded by a manual investigation involving 9 participants who examined 1,600  $\langle \text{code}, \text{summary} \rangle$  pairs. This manual analysis aimed to define a taxonomy of noisy data categories.

The taxonomy features comments- and code-related noisy data, such as *commented-out method* or *empty function*. Based on such categories, Shi *et al.* defined a set of heuristic rules and implemented them in the CAT tool. CAT works with two possible strategies based on the issues found: On the one hand, it fixes instances with minor issues. For example, it removes block-level comments. On the other hand, it completely drops instances where no fix is possible, including, for example, getters and setters.

### B. Fine-Grained Filtering: SIDE

SIDE (**S**ummary **a**lignment to **co**De **s**Emantics) is a novel quality-aware metric presented by Mastropaolo *et al.* [33]. SIDE addresses the shortcomings of traditional metrics such as BLEU [34], ROUGE [39], and METEOR [40] used in code summarization tasks.

SIDE employs a contrastive learning approach to determine the accuracy with which a code summary documents the underlying code, explicitly focusing on Java methods. Contrastive learning aims to maximize the distance between the reference code and inappropriate comments while minimizing the distance to suitable comments. The model that implements SIDE, MPNet [41], provides a continuous score ranging from -1 to 1. Scores closer to -1 indicate poor alignment between the code summary and the actual code, whereas scores closer to 1 suggest a strong alignment. SIDE showed a high correlation with human evaluations of summary quality, outperforming established metrics like BLEU, ROUGE, and METEOR.

The benefits of a quality-aware metric like SIDE extend beyond evaluating code summarization techniques, as Mastropaolo *et al.* [33] noted. SIDE can be valuable when distinguishing high-quality code documentation from subpar examples is essential. We conjecture that SIDE ensures that only the instances most likely to enhance the training procedure are used, enabling the model to converge faster without a drop in performance by filtering out low-quality elements.

Our study relies on SIDE to select  $\langle code, summary \rangle$  pairs with high coherence. Specifically, given a training set  $T$  and a threshold  $t$ , we select the training instances  $\{p_i \in T \mid \text{SIDE}(p_i) \geq t\}$ .

### III. STUDY DEFINITION, DESIGN AND PLANNING

The *goal* of this study is to empirically evaluate how code-comment coherence, through a quality-aware selection strategy grounded on SIDE, impacts the effectiveness and training efficiency of neural code summarization models.

More specifically, the study aims to address the following research questions:

**RQ<sub>0</sub>**: *How do code summarization datasets measure up in terms of code-comment coherence?* In this preliminary question, we assess the coherence of code-comment pairs of datasets commonly used in code summarization. As we aim to use a coherence-aware strategy to optimize training sets, first of all, we would like to see how the coherence is distributed.

**RQ<sub>1</sub>**: *How does a coherence-aware strategy selection impact the performance of neural code summarization models?* In this research question, we investigate how a targeted selection of training data based on code-comment coherence impacts the performance of neural code summarization models.

**RQ<sub>2</sub>**: *How does the coherence-aware strategy selection compare with a random baseline?* In this research question, we test our hypothesis that code-comment coherence is a quality attribute that can be used to select training instances.

#### A. Context Selection

The *context* of our study consists of datasets containing pairs of *Java* methods with the associated summaries. For fine-tuning the models, we consider the two most important datasets from the state of the art: TL-CodeSum [24], and Funcom [9].

The *TL-CodeSum* dataset [24] is specifically designed for the code summarization task. It consists of  $\sim 87k$  instances  $\langle code, summary \rangle$  extracted from GitHub repositories created from 2015 to 2016, and having at least 20 stars. In detail, Hu *et al.* [24] extracted the first sentence—likely to describe the overall method functionality—from the doc of each pair.

Similarly to TL-CodeSum, the *Funcom* dataset [9] is also specifically designed for code summarization. *Funcom* consists of over 2.1M  $\langle code, summary \rangle$  pairs collected from the Sourcerer repository. As for TL-CodeSum, LeClair *et al.* [9] only consider methods with their javadoc, extracting the first sentence as corresponding *summary*.

Shi *et al.* [30] found many noisy instances and duplicates in the above-described datasets and cleaned them up using their heuristic-based dataset-cleaning approach. For this reason, we use the cleaned versions of *TL-CodeSum* and *Funcom* provided by Shi *et al.* [30]. The cleaned *TL-CodeSum* contains 53,597 training instances, while the cleaned *Funcom* contains 1,184,438 training instances.

The above datasets are built automatically, and no manual check was performed, *i.e.*, there is no guarantee of their quality. For this reason, we use two additional, manually curated datasets to test the models. The first one is *CoderEval* [37], which consists of 230 Python and 230 *Java* code generation problems collected from open-source, high-starred projects which include *original* and *human-labeled* docstrings that should act as prompt for Code Generation models to generate the corresponding *code*. The instances have been subject to manual screening, for which the main criterion is the probability of appearing in real-development scenarios. We focus on the *Java* set of problems, inverting the input and the output *i.e.*, from  $\langle docstring, code \rangle$  to  $\langle code, docstring \rangle$ . To align the format of the pairs format, we performed an additional manual analysis in which one of the authors checked all the triplets with a second author to confirm the analysis. We found that some of the *docstring*(s) contained more than a sentence. Therefore, to make them consistent with the previous dataset format (*e.g.*, single sentence), we extracted the first sentence from each *docstring*. Still, we found 12 occurrences in which the corresponding *original docstring* does not describe the *code* (*e.g.*, “@inheritDoc”, “@param modelName model name of the entity”, and similar). We also excluded *docstring*: “Computes floor( $\$log\_2 (n)\$$ )  $\$+ 1\$$ .” since it includes a formula not explained in natural language. Again, to appropriately align the evaluation, we do not evaluate such instances, ending up with 218 *original* instances.

The second manually-curated dataset we use is the one by Mastropaolo *et al.* [38]. The dataset consists of 892 methods associated with their summary (*i.e.*, first sentence of the method documentation), collected from non-fork GitHub *Java* repositories with at least 300 commits, 50 contributors, and 25 stars. Such instances are in the form  $\langle summary, code \rangle$  and, as for CoderEval [37], we inverted the input and the output *i.e.*,  $\langle code, summary \rangle$ . Mastropaolo *et al.* analyzed such pairs to ensure their quality. We manually analyzed and cleaned them further (*e.g.*, “Adds an @link CarrierService to the @linkCarrier” into “Adds an CarrierService to the Carrier”), as we had done for CoderEval. No instances were removed during such a manual analysis.

We remove the instances from the test sets which appear in the training sets of *TL-CodeSum* and *Funcom*. As a result, we remove ten instances from CoderEval, which are present only in the *TL-CodeSum* training set.

#### B. Study Methodology

To answer RQ<sub>0</sub>, we use SIDE to compute the degree to which the summaries of the studied datasets document their corresponding code. We did this for each instance of the training sets included in *TL-CodeSum* and *Funcom*. To understand the coherence of the training sets, we analyze the average and the distributions of the SIDE scores of the instances.

TABLE I  
DIFFERENT SELECTIONS FOR *TL-CodeSum* AND *Funcom* TRAINING SETS.

Selection	<i>TL-CodeSum</i>	<i>Funcom</i>
Full	53,597	1,184,438
$SIDE_{0.5}$	50,073	1,080,649
$SIDE_{0.6}$	48,146	1,031,647
$SIDE_{0.7}$	44,853	952,265
$SIDE_{0.8}$	38,733	813,998
$SIDE_{0.9}$	26,258	540,170

To answer RQ<sub>1</sub>, we use the SIDE-based filter we define in Section II. We use five threshold values, *i.e.*, 0.5, 0.6, 0.7, 0.8, and 0.9. We do not use thresholds lower than 0.5 because they would result in negligible dataset reductions (lower than 10% for both), as we will observe in the results of RQ<sub>0</sub>. We report information about the different datasets in Table I. We apply each filter on the training sets of *TL-CodeSum* and *Funcom*. Such filtering leads to the definition of five new versions of both datasets.

We fine-tune a pre-trained Transformer-based model for each dataset version, *i.e.*, both the base one and its six filtered versions, producing 12 fine-tuned models. We choose to leverage the pre-trained *CodeT5+* [36] since it has been largely used for code-related tasks [42], [43], [44] and, more important, in the code summarization approaches described above. This model is built on the backbone of the well-known T5 model by Raffel *et al.* [15], yet it benefits from specific enhancements tailored for code understanding and generation tasks. During the pre-training phase, *CodeT5+* is first trained on unimodal data, which includes code and comments, employing a combination of pre-training objectives such as span-denoising [15] and Causal Language Modeling [45], [46]. Then, it is pre-trained on bi-modal data where pre-training objectives such as text-code contrastive learning, text-code matching, and text-code causal language modeling are employed. It comes with different variants: (i) *CodeT5+* 220M, (ii) *CodeT5+* 770M, (iii) *CodeT5+* 2B, (iv) *CodeT5+* 6B, and (vi) *CodeT5+* 16B. Since our experimental design would require training, validating and testing 12 models, we decided to fine-tune the *CodeT5+* variant featuring 220M trainable parameters. This choice aligns with the goal of our investigation: Rather than proposing a new code summarization technique, we aim to use a model that offers a favorable balance between size and training time while still allowing us to observe the relevant phenomenon (if present).

Considering the extensive array of our experiments, we fine-tune for 20 epochs using a batch size of 16. Additionally, we restrict the input length to 512 tokens and the output to 128 tokens, consistent with previous studies leveraging the two datasets we used [47], [48], [49]. In addition, we conduct the fine-tuning using the standard hyperparameters for *CodeT5+*, which include the AdamW optimizer [50] and a learning rate of  $2e-5$ , which is the one recommended for (Code)T5 and also used in works leveraging such models [51], [52], [53].

To prevent overfitting, we employ early stopping [54]. After each epoch, we assess the performance of the models by computing the number of correct predictions on the validation set. In line with similar research [51], [52], we implement early stopping with patience of 5 epochs and a delta of 0.01. This means that training will stop if the model’s performance does not improve by at least 0.01 for five consecutive epochs. We then select the best-performing checkpoint before early stopping. We fine-tune a *CodeT5+* model for each training set derived from the selection strategy *i.e.*, 12 (2 datasets  $\times$  six variants).

After training the models, we assess their performance on the test set dataset that, as previously explained, are the *CoderEval* [37], and the one from Mastropaolo *et al.* [38] which we refer to as the *golden sets*.

In the inference phase, we employ a beam search decoding strategy. In detail, with  $k \in \{1, 3, 5\}$ , we allow each model to generate the  $k$  most probable candidate *summaries* for the given *code*. To evaluate the generated summaries of each model, we compute the following metrics: BLEU [34], METEOR [40], and ROUGE-L [35]. **BLEU** is a metric that expresses, within a range from 0 to 1, the similarity between a generated text (candidate) and the target one (oracle). It computes the percentage of  $n$ -grams of the generated text that appear in the target, where  $n \in \{1, 2, 3, 4\}$ . **METEOR** is computed as the harmonic mean of unigram precision and recall, with the latter weighted higher than the former. It ranges from 0 to 1. **ROUGE-L** is computed as the length of the longest common subsequence (LCS) between the generated text and the target one and measures the recall by considering the proportion of the LCS relative to the length of the target text. We do not use SIDE [33] as it was employed for selecting training instances and could therefore be unnaturally biased in favor of models trained on filtered datasets. Also, we do not compute the percentage of exact matches for three reasons. First, exact matches might underestimate the actual performances of the model. Indeed, an exact match implies a correct summary, but many alternative summaries might be as correct (or even more correct, in theory) as the ones in the ground truth for the very nature of this task. Second (also related to the previous point), *CoderEval* [37] provides two summaries for each coding instance, namely *original* (*i.e.*, the docstring collected from the original source code), and *human* (*i.e.*, the docstring written from scratch by developers during the benchmark creation [37]). The model could have correctly generated only one of them, which are, by definition, both correct alternatives, thus leading to inconsistent results. Third, the dataset provided by Mastropaolo *et al.* [38] includes three different yet semantically equivalent code summaries for each *Java* method. As previously noted, each of these alternative descriptions is a valid candidate summary.

We also perform statistical hypothesis tests (Wilcoxon signed-rank test) [55] and Cliff’s delta effect size [56] to compare the distributions of the BLEU-4, METEOR, and ROUGE-L of the predictions generated by the different models trained on the filtered training sets with those of the models

TABLE II  
PERFORMANCE METRICS ON TOP-1 PREDICTIONS FOR CODEREVAL.

Dataset	Selection	#Tokens	(% Saving)	CoderEval-Original [37]			CoderEval-Human [37]			Mastropaolo <i>et al.</i> [38]		
				BLEU-4	METEOR	ROUGE	BLEU-4	METEOR	ROUGE	BLEU-4	METEOR	ROUGE
TL-CodeSum [24]	Full	↑ 7.9M	–	11.72	17.84	35.01	6.41	14.28	30.51	6.37	13.04	27.09
	SIDE <sub>0.5</sub>	7.4M	7%	12.41	17.73	35.41	6.32	14.28	31.08	6.36	12.93	27.47
	SIDE <sub>0.6</sub>	7.1M	10%	13.22	17.98	36.49	6.98	14.03	31.06	6.61	13.07	27.49
	SIDE <sub>0.7</sub>	6.6M	16%	13.17	17.93	36.14	6.79	14.83	31.90	6.30	12.95	27.60
	SIDE <sub>0.8</sub>	5.6M	28%	11.99	17.54	34.60	6.23	13.94	29.72	6.02	13.01	27.63
	SIDE <sub>0.9</sub>	↓ 3.7M	51%	11.60	17.07	33.67	6.56	14.19	30.53	5.62	12.75	27.26
Funcom [9]	Full	↑ 108.7M	–	14.25	17.61	36.53	5.93	12.95	28.36	6.77	12.94	28.00
	SIDE <sub>0.5</sub>	99.5M	9%	15.04	18.16	37.08	6.15	13.14	29.05	6.84	12.87	27.93
	SIDE <sub>0.6</sub>	95.0M	13%	16.19	19.30	38.38	7.02	14.07	30.25	7.03	12.99	28.33
	SIDE <sub>0.7</sub>	87.7M	20%	14.77	18.92	37.39	7.49	14.19	30.14	6.62	12.79	27.78
	SIDE <sub>0.8</sub>	74.2M	31%	14.10	18.34	37.04	6.53	13.38	28.45	6.93	12.78	27.99
	SIDE <sub>0.9</sub>	↓ 49.0M	54%	13.65	18.08	37.12	6.78	13.61	30.07	6.81	12.95	28.07

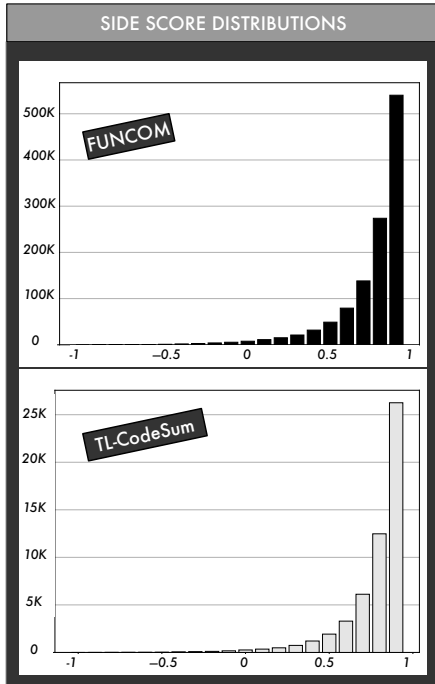


Fig. 2. Distribution of SIDE scores for *TL-CodeSum* and *Funcom* training instances.

trained on the full training sets. We use Holm’s correction [57] to adjust the  $p$ -values for the multiple tests. We reject the *null hypothesis* (there is no difference between the effectiveness of two given models) if the  $p$ -value is lower than 0.05.

Finally, we study the Pareto front to analyze the cost-benefit trade-offs between the effectiveness of the models trained on the different selections of *TL-CodeSum* and *Funcom* (benefit, measured with *i.e.*, BLEU, METEOR, and ROUGE-L) and the corresponding training dataset size (cost).

To answer RQ<sub>2</sub> we compare the selection strategy with SIDE<sub>0.9</sub> (*i.e.*, the most restrictive selection), with a *Random* baseline. In detail, we randomly sample the same number

of training instances as those selected with SIDE<sub>0.9</sub> from each dataset. We compare the effectiveness of the models trained with the training instances selected with SIDE<sub>0.9</sub> and *Random* measured in terms of the previously described metrics (*i.e.*, BLEU-4, METEOR, and ROUGE-L). Again, we perform statistical hypothesis tests (Wilcoxon signed-rank test) [55] and compute the Cliff’s delta effect size [56] to compare the distributions of BLEU-4, METEOR, and ROUGE-L of the predictions generated by the SIDE<sub>0.9</sub> model and the *Random* baseline. We use Holm’s correction [57] to adjust the  $p$ -values for the multiple tests. We reject the *null hypothesis* (there is no difference between the two models) if the  $p$ -value is lower than 0.05.

#### IV. RESULTS

In this section, we present the outcomes of our study, addressing the research questions formulated in Section III.

##### A. RQ<sub>0</sub> Consistency Analysis

Fig. 2 reports the SIDE score distributions for the training instances of the examined datasets. We can observe that most of the instances of the two datasets exhibit high-quality features, according to SIDE. In detail, the mean SIDE scores for the training sets are 0.83 and 0.81 for *TL-CodeSum* and *Funcom*, respectively. This result denotes a high coherence in the  $\langle code, summary \rangle$  pairs. Also, we can notice that, for more than 90% of the instances (93% for *TL-CodeSum* and 91% for *Funcom*), the SIDE score is greater than 0.5. While such a percentage is quite high, it is worth noting that, as a consequence, a large number of instances (3,525 and 103,789 for *TL-CodeSum* and *Funcom*, respectively) exhibit a SIDE score lower than 0.5. Besides, 51% and 54% of the instances from *TL-CodeSum* and *Funcom*, respectively, have a SIDE score lower than 0.9. While many of such instances are not necessarily detrimental to the model, they might not be beneficial either. Such distributions show a sufficiently high margin of improvement for both datasets.

## Answer to RQ<sub>0</sub>

The training sets from both *TL-CodeSum* and *Funcom* exhibit a high SIDE score, on average ( $> 0.8$  for both). Still, over half the instances have a sub-optimal SIDE score ( $< 0.9$ ). This distribution suggests that a dataset reduction approach could improve the performance of an automated code summarization model trained on it.

### B. RQ<sub>1</sub>. Selection Strategy Impact

Table II reports the results obtained from the different models fine-tuned on the full and filtered training sets, while the detailed results of the Wilcoxon signed-rank tests are in our replication package for space reasons [59]. The column “Dataset” indicates the training set from which the selections were performed, while the “Selection” column reports the SIDE threshold used to filter the training sets *i.e.*, *full*, and  $SIDE_{0.5}$  up to  $SIDE_{0.9}$ . For example, the first line in correspondence of *TL-CodeSum* (*i.e.*, Full) represents the model fine-tuned on the full training set of *TL-CodeSum*, while the one below *i.e.*, ( $SIDE_{0.5}$ ) represents the model fine-tuned on the training set filtered with SIDE score greater than 0.5. Furthermore, we report the number of tokens seen during the fine-tuning process (“Tokens” column) and the percentage of the saved training instances with respect to the complete training sets (“Saving” column). Finally, the columns “BLEU-4”, “METEOR”, and “ROUGE” report the performance in terms of percentage for the golden sets including *CoderEval-Original*, *CoderEval-Human*, and the one by Mastropaolo *et al.* [38].

In the following, we discuss results achieved through the SIDE-based selection by comparing them against the *CoderEval-Original* ground truth. Looking at the rows corresponding to *TL-CodeSum*, we observe that the model trained with the complete training set achieves 11.72% BLEU-4, 17.84% METEOR, and 35.01% ROUGE-L. Surprisingly, even though such a model underwent fine-tuning with the full training set, it is not the one exhibiting the best results in terms of these metrics. Instead, the model fine-tuned with the selection strategy with  $SIDE_{0.6}$  is the one that achieves the best results for all metrics despite being trained with  $\sim 1M$  fewer tokens during training than the one leveraging the whole corpus of code tokens. If we look at the “minimal” selection strategy *i.e.*,  $SIDE_{0.9}$ , we observe a limited drop in performance ( $\downarrow 0.12$ ,  $\downarrow 0.77$ , and  $\downarrow 1.34$ ) while saving the 51% of the training instances.

Looking at the rows corresponding to *Funcom*, we observe the same trend as for *TL-CodeSum*. The model fine-tuned on the whole training set (108.7M tokens) performs worse than the models fine-tuned with less training data. We can notice a negligible drop in performance for  $SIDE_{0.9}$  (49.0M tokens) on BLEU-4 metric ( $\downarrow 0.60$ ), while achieving slightly better performance on METEOR ( $\uparrow 0.47$ ) and ROUGE-L ( $\uparrow 0.59$ ).

The Wilcoxon signed-rank tests indicate that there is no statistically significant difference between the results obtained

by the model fine-tuned on the full training set and those fine-tuned on the filtered training sets (minimum  $p$ -value  $> 0.2$  for *TL-CodeSum*, and  $p$ -value  $> 0.08$  for *Funcom*). Also, the Cliff’s  $\delta$  effect size is always negligible for the same comparisons. Such results occur despite the most aggressive selection only leveraging 46% of the original training instances. In other words, there is little to lose in terms of summary quality when the model training set is heavily optimized.

Similar observations apply to the *CoderEval-Human* dataset. With no statistically significant differences, (minimum  $p$ -value  $> 0.23$  for *TL-CodeSum*, and  $p$ -value  $> 0.15$  for *Funcom*), according to the Wilcoxon signed-rank test and a negligible Cliff’s  $\delta$  effect size, the models specialized to produce meaningful code summaries with the filtered dataset have comparable performance to those fine-tuned with the original dataset.

For example, the model instructed with “high-quality” examples included in the  $SIDE_{0.9}$  dataset — even if trained with 4.2M fewer tokens — obtains a slightly higher BLEU-4 ( $\uparrow 0.15$ ) and a slightly lower METEOR ( $\downarrow 0.09$ ) and ROUGE-L ( $\uparrow 0.02$ ). We also observed no substantial differences in the context of *Funcom*. In this case, however, the models fine-tuned with the filtered training sets perform slightly better than the model fine-tuned on the full-training set in terms of all the metrics ( $\uparrow 0.85$  BLEU-4,  $\uparrow 0.66$  METEOR, and  $\uparrow 1.71$  ROUGE-L).

Finally, the conclusions above (no statistically significant differences, negligible Cliff’s  $\delta$  effect sizes) are further confirmed on the dataset by Mastropaolo *et al.* [38]. The model that underwent fine-tuning with the full-sized *TL-CodeSum* training set performs only slightly better than the models fine-tuned on the filtered datasets. At the same time, it is interesting to notice how—although differences are still not statistically significant—the model fine-tuned on the filtered *Funcom* training set with  $SIDE_{0.6}$  exhibits marginally better performance compared to the model that was exposed to the largest amount of code tokens, *i.e.*, the full-sized *Funcom* dataset, during training. This suggests that there may be contexts for which the selection based on the SIDE metric helps to prune out instances worsening the summarization quality.

Fig. 3 depicts the Pareto fronts for the models’ performance and training dataset size across the different selections of the *TL-CodeSum* and *Funcom* datasets (*i.e.*, Full,  $SIDE_{0.5}$ ,  $SIDE_{0.6}$ ,  $SIDE_{0.7}$ ,  $SIDE_{0.8}$ ,  $SIDE_{0.9}$ ), evaluated on the three benchmarks *i.e.*, *CoderEval-Original*, *CoderEval-Human*, and Mastropaolo *et al.*

For the models trained on the *TL-CodeSum* selections, we observe that as the training dataset size increases from  $SIDE_{0.9}$  with 3.7M tokens to  $SIDE_{0.7}$  with 6.6M tokens, the ROUGE scores remain almost constant. A slight improvement in ROUGE scores can be observed at 7.1M tokens ( $SIDE_{0.6}$ ) on *CoderEval-Original* and *CoderEval-Human*. However, at larger sizes with 7.4M tokens ( $SIDE_{0.5}$ ) and the 7.9M tokens (Full) the scores decline. At the same time, the METEOR and BLEU-4 scores remain almost constant across all dataset

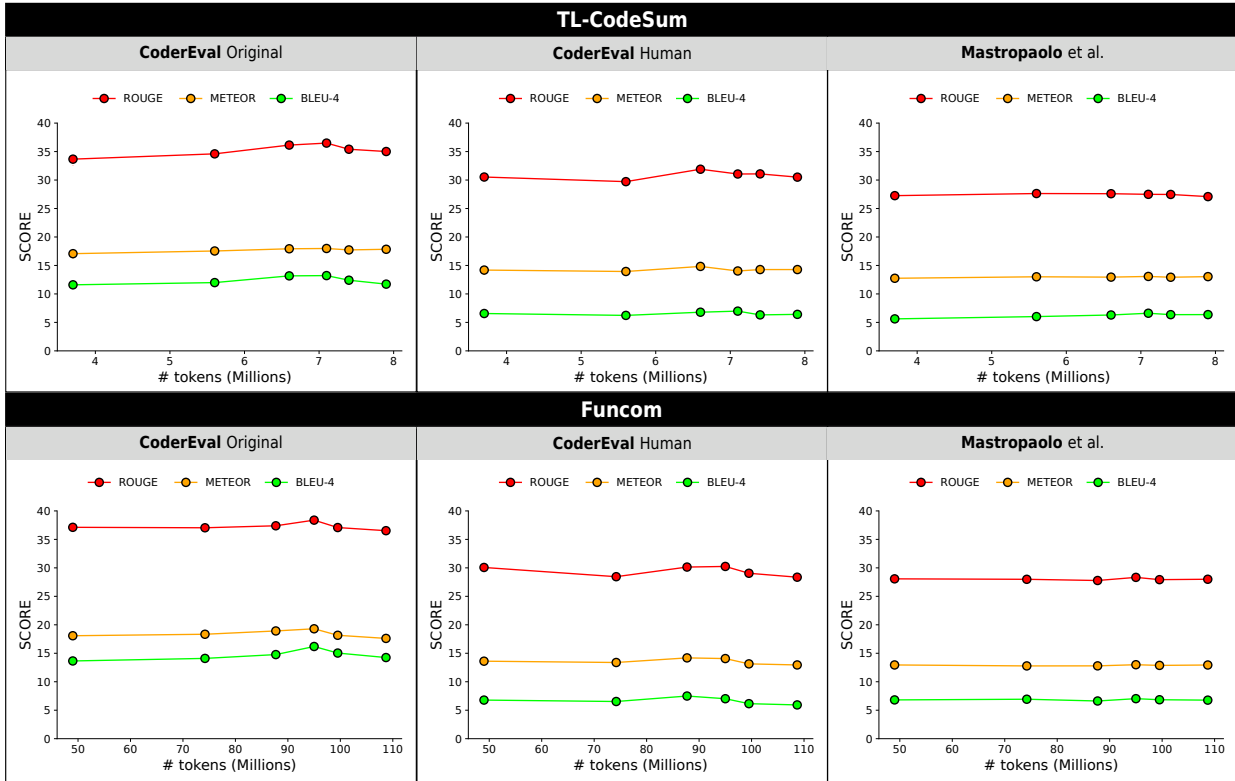


Fig. 3. Pareto front for TL-CodeSum [24] (top) and Funcom [58] (bottom). From left to right, the chart shows the number of tokens for  $SIDE_{0.9}$ ,  $SIDE_{0.8}$ ,  $SIDE_{0.7}$ ,  $SIDE_{0.6}$ ,  $SIDE_{0.5}$ , and the full dataset.

sizes, showing minimal improvements. In the Mastropaolo *et al.* dataset, the three metrics remain constant across all filtering levels, evidencing that increasing dataset size does not meaningfully impact the models’ effectiveness.

The models trained on *Funcom* selections exhibit similar results. As the dataset size decreases from the *Full* set of 108.7M tokens through *SIDE* thresholds down to  $SIDE_{0.9}$  with 49.0M, the ROUGE scores exhibit slight improvement up to  $SIDE_{0.6}$  but remain relatively stable across the golden sets, while METEOR and BLEU-4 scores follow the trend of the ROUGE scores although in a less noticeable way. For *CoderEval Human*, ROUGE scores exhibit minor improvements, while METEOR and BLEU-4 scores show no significant changes. For the Mastropaolo *et al.* dataset, the effectiveness remains stable across the three metrics and training sizes, further underscoring that increasing the training dataset size has a small impact on performance for high-quality test sets.

Although we observe slight improvements with larger dataset selections (specifically at  $SIDE_{0.6}$  and  $SIDE_{0.7}$ ), the gains are minimal, resulting in a generally flat trend on the Pareto front. Finally, as shown in Fig. 3, increasing the training dataset size does not lead to substantial improvements in downstream performance. Based on our results, we can state that increasing dataset size provides negligible improvements, with  $SIDE_{0.9}$  achieving an optimal balance between performance and resource efficiency.

### Answer to RQ<sub>1</sub>

Fine-tuning neural code summarization models with coherent code-comment instances selected through *SIDE* leads to performances comparable to those obtained when fine-tuning the model on the complete training sets, with a reduction of up to 50% of the training instances.

### C. RQ<sub>2</sub>. Comparison with the Random Baseline

Table III reports the performance achieved by the models trained on the  $SIDE_{0.9}$  and *Random* selections for both datasets (*i.e.*, *TL-CodeSum* and *Funcom*). We report such information as described in Section IV-B. Instead, Table IV reports the results of the Wilcoxon signed-rank tests with the adjusted *p*-values (column *p-value*), and the corresponding Cliff’s  $\delta$  effect sizes (column Cliff’s  $\delta$ ) for each golden set, including *CoderEval-Original*, *CoderEval-Human*, and Mastropaolo *et al.*

The results are quite comparable for the models trained on the *TL-CodeSum* dataset selections. In detail, on *CoderEval-Original*, the model trained on the  $SIDE_{0.9}$  selection achieves slightly better results ( $\uparrow 0.83$ ,  $\uparrow 0.71$ ,  $\uparrow 1.0$ ), even though without statistically significant statistical differences (*p*-value = 1.0), and negligible effect sizes for all metrics ( $\sim 0.00$ ). We obtain consistent results for *CoderEval-Human*. Instead, we

TABLE III  
PERFORMANCE METRICS ON TOP-1 PREDICTIONS FOR  $SIDE_{0.9}$  AND *Random*.

Dataset	Selection	CoderEval-Original [37]			CoderEval-Human [37]			Mastropaolo <i>et al.</i> [38]		
		BLEU-4	METEOR	ROUGE	BLEU-4	METEOR	ROUGE	BLEU-4	METEOR	ROUGE
<i>TL-CodeSum</i> [24]	$SIDE_{0.9}$	11.60	17.07	33.67	6.56	14.19	30.53	5.62	12.75	27.26
	<i>Random</i>	10.77	16.36	32.67	6.12	13.44	28.78	5.85	12.89	27.13
<i>Funcom</i> [9]	$SIDE_{0.9}$	13.65	18.08	37.12	6.78	13.61	30.07	6.81	12.95	28.07
	<i>Random</i>	13.67	17.95	36.09	6.88	13.95	29.90	5.84	12.32	27.12

TABLE IV  
RESULTS COMPARISON FOR CODESUM AND FUNCOM

Dataset	Metric	CoderEval-Original [37]		CoderEval-Human [37]		Mastropaolo <i>et al.</i> [38]	
		<i>p</i> -value	Cliff $\delta$	<i>p</i> -value	Cliff $\delta$	<i>p</i> -value	Cliff $\delta$
<i>TL-CodeSum</i> [24]	BLEU-4	1.0	0.03	0.5	0.03	1.0	0.02
	METEOR	1.0	0.02	0.5	0.014	0.4	-0.013
	ROUGE-L	1.0	0.03	0.5	0.052	1.0	0.01
<i>Funcom</i> [9]	BLEU-4	1.0	-0.004	1.0	-0.021	0.017	0.042
	METEOR	1.0	0.01	1.0	0.014	0.11	0.025
	ROUGE-L	1.0	0.03	1.0	0.013	0.07	0.021

observe slight differences in performance for the Mastropaolo *et al.* golden set. Specifically, the *Random* selection allows to achieve better BLEU-4 ( $\uparrow$  0.23) and METEOR ( $\uparrow$  0.14), but lower ROUGE ( $\downarrow$  0.13). Again, we observe no statistically significant differences and negligible effect sizes.

What was observed above is also generally true for the models trained on the *Funcom*. The model trained on the  $SIDE_{0.9}$  selection, achieves better ROUGE ( $\uparrow$  1.03,  $\uparrow$  0.17,  $\uparrow$  0.95) on all the three golden sets. Conversely, for the BLEU and METEOR metrics, the models show alternate better performance, except on the Mastropaolo *et al.* golden set. Table IV reports a statistically significant difference between the performance of the two models for the BLEU-4 metric, although with negligible effect sizes.

### Answer to RQ<sub>2</sub>

Filtering training instances based on code-comment coherence provides models with comparable effectiveness to those trained on randomly selected instances.

## V. DISCUSSIONS

Our findings challenge the conjecture that code-comment coherence, as measured by  $SIDE$  [33], is a critical quality attribute for filtering instances of code summarization datasets. By selecting  $\langle code, summary \rangle$  pairs with high-coherence for training allow to achieve the same results that would be achieved by randomly selecting such a number of instances. At the same time, we observed that reducing the datasets size up to 50% of the training instances does not significantly affect the effectiveness of the models, even when the instances are randomly selected. These results have several implications.

First, code-comment consistency might not be a problem in state-of-the-art datasets in the first place, as also suggested in the results of RQ<sub>0</sub>. Also, the DL models we adopted (and, probably, bigger models as well) are not affected by incon-

sistent code-comment pairs, even when these inconsistencies are present in the training set. Despite the theoretical benefits of filtering by  $SIDE$  [33], that is the state-of-the-art metric for measuring code-comment alignment, our results indicate its limitations in improving the *overall* quality of the training sets for code summarization task. Nevertheless, other quality aspects of code and comments that have not been explored yet (such as readability) may be important for smartly selecting the training instances. Future work should explore such quality aspects further.

Our results clearly indicate that state-of-the-art datasets contain instances that do not contribute to improving the models' effectiveness. This finding is related to a general phenomenon observed in Machine Learning and Deep Learning. Models reach convergence when they are trained for a certain amount of time (epochs). Additional training provides smaller improvements and increases the risk of overfitting. We show that the same is true for data. In terms of effectiveness, model convergence is achieved with fewer training instances than previously assumed. Limiting the number of epochs may make it possible to reach model convergence with a subset of training data, maintaining model effectiveness, reducing resource demands and minimizing the risk of overfitting. Future work could explore different criteria for data selection that identify the most informative subsets for training. Conversely, this insight suggests that currently available datasets suffer from poor diversity (thus causing the previously discussed phenomenon). This latter insight constitutes a clear warning for researchers interested in building code summarization datasets, which should include instances that add relevant information instead of adding more data, which might turn out to be useless.

Finally, it is worth pointing out that another benefit of the reduction we performed is the environmental impact. Reducing the number of training instances implies a reduced training time, which, in turn, lowers the resources necessary to perform training and, thus, energy consumption and CO<sub>2</sub> emissions. We performed a rough estimation of the training time across different selections of *TL-CodeSum* and *Funcom* datasets and estimated a proxy of the CO<sub>2</sub> emissions for each model training phase by relying on the ML CO<sub>2</sub> impact calculator<sup>1</sup> [60]. Such a calculator considers factors such as the total training time, the infrastructure used, the carbon efficiency, and the amount of carbon offset purchased. The estimation

<sup>1</sup><https://mlco2.github.io/impact/#compute>

of CO<sub>2</sub> emissions needed to train the model with the *Full* selection of *Funcom* ( $\sim 200$  hours) is equal to 26.05 Kg, while with the optimized training set, *i.e.*, *SIDE*<sub>0.9</sub> ( $\sim 90$  hours), the estimation is 11.69 Kg of CO<sub>2</sub> (-55% emissions). While we recognize that this method provides an estimation rather than a precise measurement, it offers a glimpse into the environmental impact of applying data reduction.

## VI. THREATS TO VALIDITY

This section describes the threats that could affect the findings of our study.

**Construct validity** threats concern the relationship between theory and observation. A first threat is related to how we assess the quality of the generated summaries. We leverage metrics widely used in literature and, specifically, BLEU-4, METEOR, and ROUGE. We are aware that such metrics may not fully reflect the developers’ perception of a summary quality.

Another threat could be related to the choice of *SIDE* as a driver for selecting a training set based on summary coherence. As shown in previous work [33], this metric correlates with human-based summary evaluation better than other state-of-the-art metrics.

**Internal validity** threats concern factors internal to our study that could affect our findings. One factor is related to the hyperparameter calibration of the performed training. As explained in Section III, our choices are based on those of previous studies [51], [52].

Another factor is the choice of the *SIDE* cut thresholds. We have chosen five thresholds varying from 0.5 to 0.9 and reported the findings for such levels of dataset filtering. While we did not explore the full range of *SIDE* values, we have evaluated, with a discretization of 0.1, the whole range from 0.5 above. We conjecture that lower values would not show results much different from the full datasets because the 0.5 threshold indicates a very small reduction (7%-9% on the two datasets).

Finally, we are aware that a more accurate computation of training times would require multiple runs. However, this was unfeasible given the number of configurations to evaluate and the training time required for each of them.

**Conclusion validity** threats concern the relationship between the experimentation and outcome. The study is mostly observational. Therefore, we report and discuss results through descriptive statistics. However, wherever appropriate—and in particular for RQ<sub>1</sub> and RQ<sub>2</sub>, we complement them with suitable statistical procedures (Wilcoxon rank-sum test and Cliff’s *d* effect size). Also, since multiple tests have been performed, we adjust the *p*-values using Holm’s correction procedure [57].

**External validity** threats concern the generalizability of our findings. We have considered two datasets for the training. These datasets have been specifically designed for code summarization and cleaned up using the CAT approach by Shi *et al.* [30]. As for the test set, we have considered two datasets that (i) do not overlap with the training set and (ii)

have been evaluated by humans. Although these datasets are particularly suitable for the study reported in this paper, we cannot exclude that applying the selection strategy to other datasets, particularly those related to different programming languages or application domains, might lead to different results.

## VII. RELATED WORK

In this section, we examine key literature concerning the impact of data quality on neural models for automating software engineering practices, with a particular focus on code summarization. Following this, we provide a section covering the main approaches in code summarization literature.

### A. The Importance of Data Quality for DL-Based Methods in Code Summarization

Data quality is essential for the success of DL methods designed to automate software engineering tasks. Models trained on noisy data frequently suffer significant performance degradation when deployed in real-world settings. A recent study by Shi *et al.* [61] examines various factors that could impact the evaluation of code summarizers, including metrics, code preprocessing operations, and datasets. Their empirical study showed that these elements significantly influence the evaluation of code summarization models. Specifically, Shi *et al.* observed that certain code preprocessing operations (*e.g.*, converting all tokens to lowercase) could substantially affect performance. This can either enhance the model’s effectiveness or lead to a decline in performance, showcasing the fundamental role of preprocessing strategies to either boost or impair model outcomes.

LeClair and McMillan [58] also explored the effects of different preprocessing decisions on datasets used for code summarization. Specifically, they analyzed how splitting training and test datasets by function versus by project influences model performance.

Gros *et al.* [62] investigate the relationship between code comments and code itself, particularly exploring the premise that generating code comments is akin to translating between natural languages. This conceptual alignment has facilitated using models and evaluation metrics from Natural Language Processing (NLP) in tasks like code summarization. Gros *et al.* assessed how code-comment datasets, compare with datasets used for natural language translation. Specifically, they contrasted each code-comment dataset against WMT19 [63], a renowned corpus utilized for training natural language translators. The outcomes of their research indicated that code comments are more repetitive than English sentences found in natural language, a characteristic that influences performance evaluation metrics.

Sun *et al.* [64] illustrate the issue of poor natural language data quality within the context of code search—specifically, retrieving the “closest” code matching a natural language query in a codebase. By implementing a semantic query cleaning module for code search datasets, Sun *et al.* showed that their filtering framework not only enhances the accuracy of models

trained on the filtered dataset but also conserves computing resources. Along these lines, Li *et al.* [65] investigated the quality of commit messages and found that (i) the quality of commit messages plays a fundamental role when it comes to software defects, and (ii) the overall quality of these messages declines over time, despite the developers' belief who think their commit messages are improving.

Xu *et al.* [66] explored the impact of data quality on just-in-time obsolete comment detection. They empirically showed that applying a set of manually derived rules could enhance accuracy by up to 10.7%

Shi *et al.* [30] implemented a similar approach to address data quality issues in code summarization benchmarks by introducing CAT (Code-comment Cleaning Tool). This tool can identify noisy data in various programming languages, including Java and Python. Specifically, the rule-based method (*i.e.*, CAT) developed by Shi *et al.* is designed to detect specific patterns of noisy data at both the comment and code levels, based on a taxonomy of data preprocessing noises identified across four popular datasets. An initial evaluation with CAT showed that commonly used code summarization benchmark datasets for Python and Java include noisy data pairs  $\langle \text{code}, \text{comment} \rangle$ , ranging from 31% to 66%. These noisy elements were either "fixed" to decrease their noise level or completely removed. Following this cleanup, state-of-the-art neural code summarization techniques were retrained from the ground on these refined benchmarks. The performances of the different models were then compared against the cleaned test datasets processed using CAT. This comparison indicated that the optimized training dataset significantly improved summarization accuracy and overall performance.

### B. Automated Code Summarization

Different studies have explored the automation of code summarization, with three primary approaches emerging in the literature: Information Retrieval (IR), Deep Learning (DL), and hybrid methods combining both techniques.

One of the first approaches to DL-based code summarization is the work by Iyer *et al.* [67], who introduced an RNN-based model with an attention mechanism for generating code summaries. They used an encoder-decoder architecture, proposing CODE-NN to generate summaries. It was trained on code-description pairs from StackOverflow and it demonstrated substantial improvements over traditional approaches in generating summaries, highlighting one of the first evidence of DL effectiveness for code summarization tasks.

Zhang *et al.* [10] proposed a retrieval-based neural model, Rencos, which combines information retrieval (IR) techniques with neural machine translation (NMT) specifically for code summarization. They first train an encoder-decoder model. Then, during the testing phase, they retrieve the most similar snippets from the training set in terms of syntax and semantics and encode them with the input. Finally, after fusing them it predicts the summary.

One of the first works exploring Transformer-based approaches for code summarization is the one by Ahmad *et al.*

[68]. In their work, they leverage the self-attention mechanism within Transformers to model the complex, long-distance dependencies in source code, aiming to generate appropriate summaries. Their results demonstrated that Transformers are effective for code summarization tasks.

With the advent of Large Language Models (LLMs), substantial progress has been made in automating code summarization. LLMs are capable of few-shot learning *i.e.*, providing task-specific examples in the prompt, allowing the model to perform the requested task. As first steps for LLM-based code summarization, Ahmed *et al.* [42] introduced the Automatic Semantic Augmentation of Prompts (ASAP) technique. This latter enhances LLMs performance by adding structured semantic information in the prompts. ASAP inserts (i) repository context, (ii) tagged identifiers, and (iii) data flow graphs directly into the prompt, guiding the model toward a wider knowledge of the code structure and functionality. Such technique allowed models like Code-Davinci and GPT-3.5 to achieve state-of-the-art performance on multiple programming languages for the code summarization task.

As part of advancing LLM-based code summarization, Sun *et al.* [69] conducted a study examining various prompting strategies and model configurations for improving performance in code summarization tasks. They evaluated prompting techniques such as zero-shot, few-shot, and chain-of-thought, finding that simpler approaches often performed as well as more complex methods.

### C. Code-comment Coherence

Detecting code comments that result to be inconsistent with respect to the given code has been the subject of several works. Tan *et al.* [70] proposed *iComment*, a tool that leverages NLP properties, machine learning models, and program analysis techniques to extract rules for a decision tree classifier able to detect potential inconsistencies between comments and code. They evaluated their tool for large-scale projects such as Linux, Mozilla, Apache, and Wine, finding several inconsistencies, including bugs and bad comments.

Liu *et al.* [71], instead, focused on detecting outdated comments during code changes. The authors proposed a machine learning-based approach that leverages 64 features related to code, comments, and their relationship before and after code changes to identify outdated comments. Using a random forest classifier, they evaluated their approach on several open-source projects and demonstrated its effectiveness in detecting outdated comments, achieving high precision and recall rates.

Wang *et al.* [72] proposed DComment, a framework to evaluate the quality of comments in source code. By analyzing both the code and its associated comments, DComment identifies meaningful patterns and relationships to determine whether a comment is coherent and relevant. DComment showed a strong ability to generalize across different projects on multiple datasets.

Xu *et al.* [73] proposed MCCL, which combines method comment detection and confidence learning to identify and

mitigate inconsistencies. The proposed approach uses advanced encoding techniques, including multi-head attention and graph neural networks, to analyze the relationships between code and comments. Additionally, a denoising component addresses noisy data and labeling errors, improving the detection process. MCCL was evaluated on over 1,500 open-source projects, demonstrating superior performance compared to existing methods, significantly improving precision, recall, and F1-score.

In this work, we adopt SIDE [33] as an approach for measuring code-comment coherence instead of the previously-mentioned approaches because, differently from them, (i) it has been shown to highly correlate with human judgment, and (ii) it allows us to filter instances based on a threshold (multiple selectivity levels).

### VIII. CONCLUSION

We presented an empirical investigation in which we studied how filtering out incoherent code-comment pairs in code summarization datasets affects the DL models trained to tackle this task in terms of effectiveness and removed training instances.

On the one hand, our results show that reducing the number of instances in the training set using a coherence-based approach driven by the SIDE metric [33] does not significantly impact the models' effectiveness, even when applying the strictest filter that reduces the training set of over 50%. On the other hand, we found that randomly selecting the same number of instances instead of selecting them based on their code-comment coherence allows us to achieve comparable results.

Our future research agenda includes investigating other quality aspects for selecting training data, such as instance diversity and readability. Furthermore, we plan to investigate how far we can go in randomly removing instances without affecting the quality of generated code summaries.

### IX. DATA AVAILABILITY

The study dataset and scripts used for the analysis are available and documented in our online replication package [59].

### X. ACKNOWLEDGMENTS

This publication is part of the project PNRR-NGEU which has received funding from the MUR – DM 118/2023. This work has been partially supported by the European Union - NextGenerationEU through the Italian Ministry of University and Research, Projects PRIN 2022 “QualAI: Continuous Quality Improvement of AI-based Systems”, grant n. 2022B3BP5S, CUP: H53D23003510006.

### REFERENCES

- [1] B. Fluri, M. Wursch, and H. C. Gall, “Do code and comments co-evolve? on the relation between source code and comment changes,” in *14th Working conference on reverse engineering (WCRE 2007)*. IEEE, 2007, pp. 70–79.
- [2] B. Fluri, M. Wursch, E. Giger, and H. C. Gall, “Analyzing the co-evolution of comments and source code,” *Software Quality Journal*, vol. 17, pp. 367–394, 2009.
- [3] M. Linares-Vásquez, B. Li, C. Vendome, and D. Poshyvanyk, “How do developers document database usages in source code?(n),” in *2015 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 2015, pp. 36–41.

- [4] F. Wen, C. Nagy, G. Bavota, and M. Lanza, “A large-scale empirical study on code-comment inconsistencies,” in *2019 IEEE/ACM 27th International Conference on Program Comprehension (ICPC)*. IEEE, 2019, pp. 53–64.
- [5] P. Rani, A. Blasi, N. Stulova, S. Panichella, A. Gorla, and O. Nierstrasz, “A decade of code comment quality assessment: A systematic literature review,” *Journal of Systems and Software*, vol. 195, p. 111515, 2023.
- [6] S. C. B. de Souza, N. Anquetil, and K. M. de Oliveira, “A study of the documentation essential to software maintenance,” in *Proceedings of the 23rd annual international conference on Design of communication: documenting & designing for pervasive information*, 2005, pp. 68–75.
- [7] S. Haque, A. LeClair, L. Wu, and C. McMillan, “Improved automatic summarization of subroutines via attention to file context,” in *Proceedings of the 17th International Conference on Mining Software Repositories*, 2020, pp. 300–310.
- [8] X. Hu, G. Li, X. Xia, D. Lo, and Z. Jin, “Deep code comment generation with hybrid lexical and syntactical information,” *Empirical Software Engineering*, vol. 25, pp. 2179–2217, 2020.
- [9] A. LeClair, S. Jiang, and C. McMillan, “A neural model for generating natural language summaries of program subroutines,” in *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 2019, pp. 795–806.
- [10] J. Zhang, X. Wang, H. Zhang, H. Sun, and X. Liu, “Retrieval-based neural source code summarization,” in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering (ICSE 2020)*, 2020, pp. 1385–1397.
- [11] A. Mastropaolo, S. Scalabrino, N. Cooper, D. N. Palacio, D. Poshyvanyk, R. Oliveto, and G. Bavota, “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, 2021, pp. 336–347.
- [12] S. Gao, C. Gao, Y. He, J. Zeng, L. Nie, X. Xia, and M. Lyu, “Code structure-guided transformer for source code summarization,” *ACM Transactions on Software Engineering and Methodology*, vol. 32, no. 1, pp. 1–32, 2023.
- [13] A. Vaswani, “Attention is all you need,” *Advances in Neural Information Processing Systems*, 2017.
- [14] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [15] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, “Exploring the limits of transfer learning with a unified text-to-text transformer,” *Journal of machine learning research*, vol. 21, no. 140, pp. 1–67, 2020.
- [16] B. Wei, G. Li, X. Xia, Z. Fu, and Z. Jin, “Code generation as a dual task of code summarization,” *Advances in neural information processing systems*, vol. 32, 2019.
- [17] A. Svyatkovskiy, S. K. Deng, S. Fu, and N. Sundaresan, “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*, 2020, pp. 1433–1443.
- [18] J. Liu, C. S. Xia, Y. Wang, and L. Zhang, “Is your code generated by ChatGPT really correct? rigorous evaluation of large language models for code generation,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [19] S. Ugare, T. Suresh, H. Kang, S. Misailovic, and G. Singh, “Improving llm code generation with grammar augmentation,” *arXiv preprint arXiv:2403.01632*, 2024.
- [20] M. Jin, S. Shahriar, M. Tufano, X. Shi, S. Lu, N. Sundaresan, and A. Svyatkovskiy, “Inferfix: End-to-end program repair with llms,” in *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2023, pp. 1646–1656.
- [21] Z. Chen, S. Kommrusch, and M. Monperrus, “Neural transfer learning for repairing security vulnerabilities in C code,” *IEEE Transactions on Software Engineering*, vol. 49, no. 1, pp. 147–165, 2022.
- [22] Z. Chen, S. Kommrusch, M. Tufano, L.-N. Pouchet, D. Poshyvanyk, and M. Monperrus, “Sequencer: Sequence-to-sequence learning for end-to-end program repair,” *IEEE Transactions on Software Engineering*, vol. 47, no. 9, pp. 1943–1959, 2019.

- [23] M. Tufano, C. Watson, G. Bavota, M. Di Penta, M. White, and D. Poshyvanyk, "An empirical investigation into learning bug-fixing patches in the wild via neural machine translation," in *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, 2018, pp. 832–837.
- [24] X. Hu, G. Li, X. Xia, D. Lo, S. Lu, and Z. Jin, "Summarizing source code with transferred API knowledge," in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*. ijcai.org, 2018, pp. 2269–2275.
- [25] X. Hu, G. Li, X. Xia, D. Lo, and Z. Jin, "Deep code comment generation," in *Proceedings of the 26th conference on program comprehension (ICPC 2010)*, 2018, pp. 200–210.
- [26] A. LeClair, S. Haque, L. Wu, and C. McMillan, "Improved code summarization via a graph neural network," in *Proceedings of the 28th international conference on program comprehension (ICPC 2020)*, 2020, pp. 184–195.
- [27] C. Bird, A. Bachmann, E. Aune, J. Duffy, A. Bernstein, V. Filkov, and P. T. Devanbu, "Fair and balanced?: bias in bug-fix datasets," in *Proceedings of the 7th joint meeting of the European Software Engineering Conference and the ACM SIGSOFT International Symposium on Foundations of Software Engineering, 2009, Amsterdam, The Netherlands, August 24-28, 2009*. ACM, 2009, pp. 121–130.
- [28] K. Herzig, S. Just, and A. Zeller, "It's not a bug, it's a feature: how misclassification impacts bug prediction," in *35th International Conference on Software Engineering, ICSE '13, San Francisco, CA, USA, May 18-26, 2013*. IEEE Computer Society, 2013, pp. 392–401.
- [29] A. Bachmann, C. Bird, F. Rahman, P. T. Devanbu, and A. Bernstein, "The missing links: bugs and bug-fix commits," in *Proceedings of the 18th ACM SIGSOFT International Symposium on Foundations of Software Engineering, 2010, Santa Fe, NM, USA, November 7-11, 2010*. ACM, 2010, pp. 97–106.
- [30] L. Shi, F. Mu, X. Chen, S. Wang, J. Wang, Y. Yang, G. Li, X. Xia, and Q. Wang, "Are we building on the rock? on the importance of data preprocessing for code summarization," in *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2022, pp. 107–119.
- [31] N. Dziri, S. Milton, M. Yu, O. Zaiane, and S. Reddy, "On the origin of hallucinations in conversational models: Is it the datasets or the models?" *arXiv preprint arXiv:2204.07931*, 2022.
- [32] A. Corazza, V. Maggio, and G. Scanniello, "On the coherence between comments and implementations in source code," in *2015 41st Euromicro Conference on Software Engineering and Advanced Applications*. IEEE, 2015, pp. 76–83.
- [33] A. Mastropaolo, M. Ciniselli, M. Di Penta, and G. Bavota, "Evaluating code summarization techniques: A new metric and an empirical characterization," in *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, 2024, pp. 1–13.
- [34] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "Bleu: a method for automatic evaluation of machine translation," in *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 2002, pp. 311–318.
- [35] C.-Y. Lin, "Rouge: A package for automatic evaluation of summaries," in *Text summarization branches out*, 2004, pp. 74–81.
- [36] Y. Wang, H. Le, A. D. Gotmare, N. D. Bui, J. Li, and S. C. Hoi, "Codet5+: Open code large language models for code understanding and generation," *arXiv preprint arXiv:2305.07922*, 2023.
- [37] H. Yu, B. Shen, D. Ran, J. Zhang, Q. Zhang, Y. Ma, G. Liang, Y. Li, Q. Wang, and T. Xie, "Codereval: A benchmark of pragmatic code generation with generative pre-trained models," in *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering*, 2024, pp. 1–12.
- [38] A. Mastropaolo, L. Pascarella, E. Guglielmi, M. Ciniselli, S. Scalabrino, R. Oliveto, and G. Bavota, "On the robustness of code generation techniques: An empirical study on Github CoPilot," in *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2023, pp. 2149–2160.
- [39] C.-Y. Lin, "Rouge: A package for automatic evaluation of summaries," in *Text summarization branches out*, 2004, pp. 74–81.
- [40] S. Banerjee and A. Lavie, "METEOR: An automatic metric for MT evaluation with improved correlation with human judgments," in *Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization, ACL*, 2005, pp. 65–72.
- [41] K. Song, X. Tan, T. Qin, J. Lu, and T.-Y. Liu, "MPNet: masked and permuted pre-training for language understanding," *ArXiv*, vol. abs/2004.09297, 2020. [Online]. Available: <https://api.semanticscholar.org/CorpusID:215827489>
- [42] T. Ahmed, K. S. Pai, P. Devanbu, and E. Barr, "Automatic semantic augmentation of language model prompts (for code summarization)," in *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, 2024, pp. 1–13.
- [43] H. N. Phan, H. N. Phan, T. N. Nguyen, and N. D. Bui, "Repohyper: Better context retrieval is all you need for repository-level code completion," *arXiv preprint arXiv:2403.06095*, 2024.
- [44] G. Yang, Y. Zhou, W. Yang, T. Yue, X. Chen, and T. Chen, "How important are good method names in neural code generation? a model robustness perspective," *ACM Transactions on Software Engineering and Methodology*, vol. 33, no. 3, pp. 1–35, 2024.
- [45] S. Soltan, S. Ananthakrishnan, J. FitzGerald, R. Gupta, W. Hamza, H. Khan, C. Peris, S. Rawls, A. Rosenbaum, A. Rumshisky *et al.*, "Alex-atm 20b: Few-shot learning using a large-scale multilingual seq2seq model," *arXiv preprint arXiv:2208.01448*, 2022.
- [46] Y. Tay, M. Dehghani, V. Q. Tran, X. Garcia, J. Wei, X. Wang, H. W. Chung, S. Shakeri, D. Bahri, T. Schuster *et al.*, "UI2: Unifying language learning paradigms," *arXiv preprint arXiv:2205.05131*, 2022.
- [47] A. Mastropaolo, N. Cooper, D. N. Palacio, S. Scalabrino, D. Poshyvanyk, R. Oliveto, and G. Bavota, "Using transfer learning for code-related tasks," *IEEE Transactions on Software Engineering*, vol. 49, no. 4, pp. 1580–1598, 2022.
- [48] Y. Zhou, J. Shen, X. Zhang, W. Yang, T. Han, and T. Chen, "Automatic source code summarization with graph attention networks," *Journal of Systems and Software*, vol. 188, p. 111257, 2022.
- [49] R. Tufano, L. Pascarella, and G. Bavota, "Automating code-related tasks through transformers: The impact of pre-training," in *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2023, pp. 2425–2437.
- [50] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [51] A. Mastropaolo, M. Di Penta, and G. Bavota, "Towards automatically addressing self-admitted technical debt: How far are we?" in *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 2023, pp. 585–597.
- [52] M. Ciniselli, A. Martin-Lopez, and G. Bavota, "On the generalizability of deep learning-based code completion across programming language versions," in *Proceedings of the 32nd IEEE/ACM International Conference on Program Comprehension*, 2024, pp. 99–111.
- [53] A. Mastropaolo, V. Nardone, G. Bavota, and M. Di Penta, "How the training procedure impacts the performance of deep learning-based vulnerability patching," in *Proceedings of the 28th International Conference on Evaluation and Assessment in Software Engineering*, ser. EASE '24, 2024, p. 150–159.
- [54] L. Prechelt, "Early stopping-but when?" in *Neural Networks: Tricks of the trade*. Springer, 2002, pp. 55–69.
- [55] F. Wilcoxon, "Individual comparisons by ranking methods," in *Breakthroughs in statistics: Methodology and distribution*. Springer, 1992, pp. 196–202.
- [56] R. J. Grissom and J. J. Kim, *Effect sizes for research: A broad practical approach*. Lawrence Erlbaum Associates Publishers, 2005.
- [57] S. Holm, "A simple sequentially rejective multiple test procedure," *Scandinavian journal of statistics*, pp. 65–70, 1979.
- [58] A. LeClair and C. McMillan, "Recommendations for datasets for source code summarization," *arXiv preprint arXiv:1904.02660*, 2019.
- [59] Anonymous, "Replication Package of "Optimizing Datasets for Code Summarization: Is Code-Comment Coherence Enough?,"" Nov. 2024. [Online]. Available: <https://doi.org/10.5281/zenodo.14060451>
- [60] A. Lacoste, A. Luccioni, V. Schmidt, and T. Dandres, "Quantifying the carbon emissions of machine learning," *arXiv preprint arXiv:1910.09700*, 2019.
- [61] E. Shi, Y. Wang, L. Du, J. Chen, S. Han, H. Zhang, D. Zhang, and H. Sun, "On the evaluation of neural code summarization," in *Proceedings of the 44th international conference on software engineering*, 2022, pp. 1597–1608.
- [62] D. Gros, H. Sezhayan, P. Devanbu, and Z. Yu, "Code to comment" translation" data, metrics, baselining & evaluation," in *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering*, 2020, pp. 746–757.
- [63] L. Barrault, O. Bojar, M. R. Costa-Jussa, C. Federmann, M. Fishel, Y. Graham, B. Haddow, M. Huck, P. Koehn, S. Malmasi *et al.*, "Findings of the 2019 conference on machine translation (wmt19)." *ACL*, 2019.

- [64] Z. Sun, L. Li, Y. Liu, X. Du, and L. Li, "On the importance of building high-quality training datasets for neural code search," in *Proceedings of the 44th International Conference on Software Engineering*, 2022, pp. 1609–1620.
- [65] J. Li and I. Ahmed, "Commit message matters: Investigating impact and evolution of commit message quality," in *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2023, pp. 806–817.
- [66] S. Xu, Y. Yao, F. Xu, T. Gu, J. Xu, and X. Ma, "Data quality matters: A case study of obsolete comment detection," in *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2023, pp. 781–793.
- [67] S. Iyer, I. Konstas, A. Cheung, and L. Zettlemoyer, "Summarizing source code using a neural attention model," in *54th Annual Meeting of the Association for Computational Linguistics 2016*. Association for Computational Linguistics, 2016, pp. 2073–2083.
- [68] W. U. Ahmad, S. Chakraborty, B. Ray, and K.-W. Chang, "A transformer-based approach for source code summarization," *arXiv preprint arXiv:2005.00653*, 2020.
- [69] W. Sun, Y. Miao, Y. Li, H. Zhang, C. Fang, Y. Liu, G. Deng, Y. Liu, and Z. Chen, "Source code summarization in the era of large language models," *arXiv preprint arXiv:2407.07959*, 2024.
- [70] L. Tan, D. Yuan, G. Krishna, and Y. Zhou, "/\* icomment: Bugs or bad comments?\*", in *Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles*, 2007, pp. 145–158.
- [71] Z. Liu, H. Chen, X. Chen, X. Luo, and F. Zhou, "Automatic detection of outdated comments during code changes," in *2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC)*, vol. 1. IEEE, 2018, pp. 154–163.
- [72] D. Wang, Y. Guo, W. Dong, Z. Wang, H. Liu, and S. Li, "Deep code-comment understanding and assessment," *IEEE Access*, vol. 7, pp. 174 200–174 209, 2019.
- [73] Z. Xu, S. Guo, Y. Wang, R. Chen, H. Li, X. Li, and H. Jiang, "Code comment inconsistency detection based on confidence learning," *IEEE Transactions on Software Engineering*, 2024.