

CupCleaner: a hybrid data cleaning approach for comment updating

Qingyuan LIANG¹, Zeyu SUN^{2*}, Qihao ZHU¹, Junhao HU¹, Yifan ZHAO¹,
Yakun ZHANG¹ & Lu ZHANG^{1*}

¹Key Laboratory of High Confidence Software Technologies (Peking University), Ministry of Education,
School of Computer Science, Peking University, Beijing 100871, China

²National Key Laboratory of Space Integrated Information System, Institute of Software,
Chinese Academy of Sciences, Beijing 100190, China

Received 19 August 2024/Revised 23 June 2025/Accepted 17 October 2025/Published online 11 June 2026

Abstract Comment updating is an emerging task in software evolution that aims to automatically revise source code comments in accordance with code changes. This task plays a vital role in maintaining code-comment consistency throughout software development. Recently, deep learning-based approaches have shown great potential in addressing comment updating by learning complex patterns between code edits and corresponding comment modifications. However, the effectiveness of these learning-based approaches heavily depends on the quality of training data. Existing datasets are typically constructed by mining version histories from open-source repositories such as GitHub, where there is often a lack of quality control over comment edits. As a result, these datasets may contain noisy or inconsistent samples that hinder model learning and generalization. In this paper, we focus on cleaning existing comment updating datasets, considering both the data's characteristics in the updating scenario and their implications on the model training process. We propose a hybrid statistical approach named CupCleaner (comment updating's CLEANER) to achieve this purpose. Specifically, we combine static semantic information within data samples and dynamic loss information during the training process to clean the dataset. Experimental results demonstrate that, on the same test set, both the individual static strategy and the dynamic strategy can significantly filter out a portion of the data and enhance the performance of the model. Furthermore, employing a model ensemble approach can combine the characteristics of static and dynamic cleaning, further enhancing the performance of the model and the reliability of its output results.

Keywords software evolution, comment updating, data cleaning, software engineering, deep learning

Citation Liang Q Y, Sun Z Y, Zhu Q H, et al. CupCleaner: a hybrid data cleaning approach for comment updating. *Sci China Inf Sci*, 2026, 69(7): 172102, <https://doi.org/10.1007/s11432-024-4637-2>

1 Introduction

Software evolution is a pivotal and indispensable process in the field of software engineering (SE), as it encompasses the continuous development and improvement of software systems [1, 2]. Throughout the process of software evolution, a substantial volume of data is generated and recorded, holding substantial research value. These data can be leveraged to support various data-driven research tasks related to software evolution, such as automated code repair and code review [3–7]. Among them, comment updating has been recognized as an important and emerging research task during software development and maintenance [8, 9]. This task aims to automatically generate new code comments based on existing code comments and code changes. On one hand, automatic comment updating can save time and effort for programmers to manually write new comments. Unlike directly generating comments, comment updating can take into account more information, such as code changes and old comments. Thus, this task is more practical and relevant to the daily development scenarios of programmers. On the other hand, if comments are not updated in a timely manner, they may mislead future development activities. For example, if obsolete TODO comments are not removed, they may introduce new bugs in future development [10].

Currently, with the advancement of artificial intelligence (AI) technology (e.g., deep learning), particularly in the area of code pre-training, there is a promising potential to offer an effective solution to the comment updating task. High-quality data is a key factor in deep learning-based approaches [11–13], and higher quality data often leads to more powerful and more reliable models, while lower quality data can have the opposite effect. Data collection

* Corresponding author (email: zeyu.zys@gmail.com, zhanglucs@pku.edu.cn)

```

// the private value length in bits , zero otherwise.
public int get j ( ) { return j ; }
// subgroup factor
public big integer get j ( ) { return j ; }

```

(a) An example of noising data sample in comment updating task.

```

// @return a hashmap of all formulas
public HashMap<String, String> getFormulas() {
    return formulas; }
// @return a list of all formulas
public PactFormulaList getFormulas() {
    return formulas; }

```

(b) An example where the update intention is clear and the code changes match the comment changes.

Figure 1 (Color online) Two data examples in the comment updating task.

during the software evolution process presents challenges, as handling different versions of code and comments often makes it difficult to ensure the quality of data. Datasets related to the comment updating task are typically crawled from online repositories, such as GitHub, by extracting commit history versions. This also makes it easier for noisy data to be included in datasets for the comment updating task [14–17]. Figure 1 illustrates two data samples in the comment updating task. In Figure 1(a), the sample is considered noisy, as the comment fails to reflect the functionality of the source code (e.g., compared to the bottom code, “subgroup factor” comment is unrelated to its functionality). Additionally, the comment change does not capture the variations between code changes (e.g., compared to the original comment, the changes in the updated comment fail to reflect the transition in the code from “int” to “big integer”). Figure 1(b) presents a clear sample that code changes match the comment changes, where the change from “HashMap” to “List” is reflected in both the comments and the code. If the dataset contains a large amount of noise like that shown in Figure 1(a), the results generated by a model trained with the dataset would also be misleading. Therefore, to enhance the reliability of the models, it is crucial to improve the quality of the dataset by cleaning datasets.

While recent studies have explored data cleaning for code-related tasks, such as comment generation or code search [18–20], these approaches primarily rely on heuristic rules targeting individual code-comment pairs. Although effective to some extent in identifying clearly low-quality samples, such rule-based methods face fundamental limitations when applied to the comment updating task. First, these methods often fail to capture the semantic alignment between code changes and comment changes. Unlike static code-comment matching tasks, comment updating inherently involves reasoning over differences between old and new code, and the corresponding changes in comments. Simple surface-level rules (e.g., keyword matching or edit distance thresholds) are insufficient for identifying whether a comment change is meaningful or consistent with the associated code edit. This highlights the need for cleaning strategies that account for semantic-level relationships between code and comment evolution. Second, existing approaches overlook the rich signals available during model training. In practice, low-quality data often reveals itself through training dynamics, for example, by producing higher loss values, slower convergence, or degraded generalization. However, rule-based methods operate independently of model feedback, treating all training samples uniformly regardless of how the model learns from them. This fails to leverage the opportunity to evaluate data quality in a model-aware manner.

In this paper, we propose CupCleaner (comment updating’s CLEANER), a hybrid data cleaning approach that leverages both static and dynamic information of the data. Specifically, we take the semantics of the data as the static information (i.e., semantic-aware), while in terms of dynamic information, we consider the loss of the model during the training process (i.e., loss-aware). The static cleaning strategy aims to address the first limitation, while the dynamic cleaning strategy aims to tackle the second limitation. By incorporating these two aspects, CupCleaner aims to enhance the data cleaning process effectively. Figure 2(a) depicts the task of comment updating,

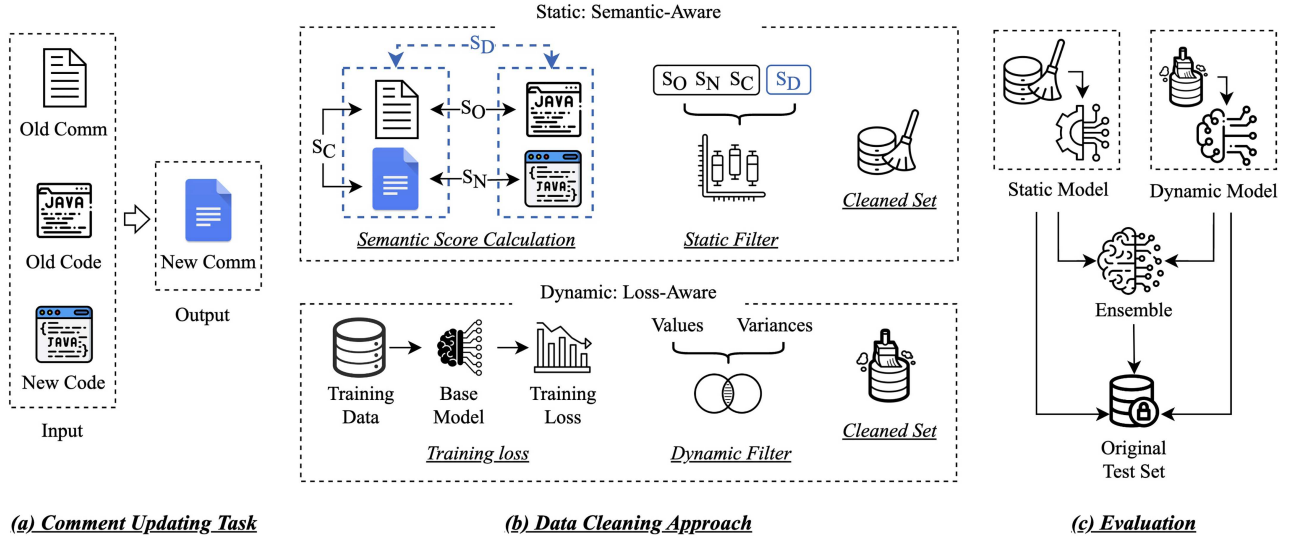


Figure 2 (Color online) Overview of our research procedure.

and Figure 2(b) shows an overview of the CupCleaner. During the process of leveraging semantic information for cleaning, we focus on two main types of noise in comment updating: weak correlation within comments or code-comment pairs (i.e., black solid line), and weak correlation between code changes and comment changes (i.e., blue dashed line). Then we design a heuristic rule based on semantic information to compute a score for each data sample in the comment updating dataset, and filter out the ones with low scores. During the data cleaning process using model training loss, we initially train a base model to derive the loss and its variations for each data sample throughout the training process. Subsequently, we analyze by considering both the values and variances of the loss, enabling us to filter out data samples that pose challenges for the model to learn effectively. As depicted in Figure 2(c), CupCleaner employs static and dynamic statistical information to filter the comment updating dataset. We train models using datasets that are cleaned separately using static and dynamic information, and conduct testing on the original test set to evaluate their performance. Additionally, we further integrate the two cleaning approaches at the model level to enhance the robustness and performance of the final model. It is worth noting that CupCleaner is an unsupervised approach for improving data quality, which allows it to work without the need for pre-prepared or assumed high-quality data.

To evaluate CupCleaner, we first conduct a human evaluation by mixing the noise data and high-quality data. The evaluation results show that, with a maximum score of 5, the average score for the noise data identified by CupCleaner is only 2.5, while the high-quality data received a score of 4.3. Then we conduct experiments on two comment updating datasets using two representative models: PLBART [21] and CodeT5 [22]. The experimental results demonstrate that combining static and dynamic cleaning approaches effectively enhances the model's performance. For example, on the dataset from Panthaplackel et al., 2021 [9], compared to models trained on original data, the PLBART model trained on data cleaned by CupCleaner increases the exact match metric from 33.33% to 38%. Similarly, the CodeT5 model trained on cleaned data also raises the exact match metric from 41.33% to 45.33%, surpassing the performance of the dedicated code editing pre-trained model (i.e., CoditT5 [23]). In addition, we separately evaluate the impact of static and dynamic cleaning strategies on model training. Experimental results indicate that both static and dynamic cleaning strategies enhance model performance. Furthermore, combining these two cleaning strategies further enhances the model's performance and generalization capabilities. Besides, we also explore the time consumption of cleaning data using CupCleaner. Comparative results indicate that the time spent cleaning a dataset is acceptable compared to fine-tuning models on the dataset. In summary, training with data cleaned using CupCleaner can provide higher-quality data within a limited time frame. This enables the same model to exhibit improved performance, thus enhancing the model's reliability.

We summarize our contributions as follows.

- **Data cleaning.** We propose semantic-aware static data cleaning strategies and loss-aware dynamic data cleaning strategies for data cleaning. To the best of our knowledge, we are the first to develop a data cleaning approach that takes into account both the characteristics of the data and its impact on the model training process.
- **Model ensemble.** We ensemble different cleaning strategies at the model level, selecting the superior output from models trained with two different cleaning strategies as the final output of the model.

- **Evaluation.** We conduct experiments on two different comment updating datasets. Human evaluation results suggest that the scores of the filtered-out data are significantly lower compared to those of the retained data. The experimental results indicate that both cleaning strategies improve data quality, and combining the two strategies further enhances the model’s performance and generalization ability.

- The data and code are available at CupCleaner¹).

2 Related work

Recently, an increasing number of researchers have become aware of the beneficial impact of high-quality data on model training [24–26]. Meanwhile, an increasing number of AI for SE studies focus on editing or updating related tasks. Below, we discuss the two most related works to our paper: code data cleaning techniques and updating-related tasks.

2.1 Code data cleaning

Data quality is crucial for deep learning models, as only high-quality datasets can comprehensively exploit the potential of deep learning techniques. In recent years, many researchers in the field of AI for SE focused on how to improve the quality of code-related data. Sun et al. [18] conducted an empirical study on the dataset for code search and found that more than one-third of the queries contain noise. To improve the quality of code search datasets, they propose a data cleaning framework by combining syntactic and semantic filters. The experimental results show that the performance of code search is significantly improved with the cleaned dataset. Shi et al. [27] focused on evaluation and data issues in code summarization tasks, and found that the pre-processing choices of the data have a significant impact on the experimental results. Huang et al. [20] analyzed the issue of whether code needs to be commented or not. They find that the percentage of methods with both header and internal comments in software systems is low. Then they design an approach to determining whether code needs to be commented by considering structural features, syntactic features, and textual features. Shi et al. [19] subsequently analyzed the noisy data in code comments and improved the quality of code comments. They propose an automatic code-comment cleaning tool and conduct experiments on four code summarization datasets. Xu et al. [28] utilized character-level differences in comment changes to clean the data used for the obsolete comment detection task. Specifically, character-level differences can be used to filter out minor comment changes. These filtered comment changes serve as non-obsolete comment samples and are used to train a classifier for determining whether comments are obsolete.

However, applying these approaches to updating-related tasks may primarily face the following two limitations. First, the approaches focusing on the correlation between individual code and comments may not align well with the updating-related scenario, as the comment updating task involves changes in both code and comments. Second, relying solely on heuristic rules or isolated semantics may not adequately capture the overall data quality. It is also crucial to consider the impact of data samples on the model training process. Different from these approaches that focus on cleaning individual comments or code, our proposed CupCleaner combines semantic-based static information and model loss-based dynamic information to consider the quality of data comprehensively. Additionally, CupCleaner is a purely unsupervised approach that does not require extra high-quality data to train a data quality discriminator. This allows it to be more flexibly adapted to various scenarios, such as being used to provide higher-quality datasets before Xu et al.’s work [28].

2.2 Comment updating tasks

Updates are a fundamental part of software development, and many intelligent software engineering tasks revolve around automatically generating or modifying artifacts during the evolution process. For example, automatic program repair focuses on updating buggy code to restore correctness [6, 29–33], while automated code review assists developers in updating code based on review suggestions [4, 34, 35].

Among these, comment updating has recently emerged as a focused task that aims to automatically revise natural language comments in response to changes in the corresponding source code [8, 9, 23, 36]. One of the earliest efforts in this direction is the work by Panthaplackel et al. [8], which formulates the task as generating an updated comment based on the code changes and its associated old comment. To support this task, the authors constructed a large-scale dataset by mining synchronized comment-code edits from commit histories of open-source Java projects. This dataset focuses primarily on `@return` comments and provides aligned (old comment, old code, new code, new comment) tuples for supervised learning. Subsequent work by the same group [9] shifts focus to the problem of

1) CupCleaner. <https://github.com/liangqingyuan/cupcleaner.hybrid>.

detecting whether a comment is consistent with its corresponding code. Although the primary task is framed as binary classification (i.e., consistent vs. inconsistent), the dataset construction process involves identifying cases where code has changed but comments remain outdated, making it possible to extract a large number of comment update samples. Compared to [8], this dataset expands the coverage beyond `@return` to include additional comment types such as `@param` and summary comments. CoditT5 [23] later reused this dataset by selecting samples where both code and comments were modified, repurposing them for the comment updating task and training a sequence-to-sequence model for generating revised comments.

However, datasets supporting such tasks are often mined from the commit versions of open-source repositories. Due to the complexity of version iterations and the diversity of update intentions, such datasets are more prone to incorporating noisy data [14–17]. To evaluate and enhance the quality of the comment updating task, we propose CupCleaner for data cleaning for comment updating. Unlike studies that focus on designing specialized deep learning approaches to enhance model generation capabilities, this paper aims to clean existing datasets to improve the performance and reliability of the original models. Our experimental results indicate that CupCleaner can effectively filter out noise from the comment updating datasets, and models trained on the cleaned data exhibit better performance.

3 Approach

To measure data quality and clean comment updating datasets, we consider the real requirements in comment updating and design the CupCleaner to automate the data cleaning process. The overall framework of our approach is shown in Figure 2(b). Our data cleaning process consists of two strategies: the static data cleaning strategy, which utilizes semantic-based heuristic rules to filter the data, and the dynamic data cleaning strategy, which leverages loss information from the model training process to filter the data.

3.1 Static strategy

In the static strategy, we design a semantic-aware criterion to calculate the quality scores for each data sample.

3.1.1 Semantic score calculation

Different from the data cleaning process that solely caters to individual comment-code pairs, samples of the comment updating task require considering the correlation not only between code and comments but also between code changes and comment changes. Therefore, we categorize the data noise into two types. The first type of noise exists between comments or comment-code pairs within each data sample. We evaluate their correlation by calculating their internal semantic similarity. The second type of weak correlation exists between comment changes and code changes, and we evaluate this correlation based on *diff* information. Below, we discuss the details of how to calculate the score and build filters.

Noise existing between comments or comment-code pairs (Type I). This type of noise arises from common issues with data quality, such as invalid or incorrect comments. We leverage the inherent semantic similarity within them to identify and address this type of noise. To determine whether there is invalid data within the comments or code, we use the code pre-trained models to convert them into embedding vectors to represent semantics. Currently, there are many pre-trained models used to represent the semantics of texts [37–41]. We choose the state-of-the-art model pre-trained on code-related datasets, UniXcoder [41], to calculate the semantics of comments or code. We use the cosine similarity of semantic embeddings to describe the correlation between old and new comments, and between comments and code, respectively. Below, we present the specific calculation details.

$$S_C^i = \text{Cos}(M(\text{comm}_{\text{old}}^i), M(\text{comm}_{\text{new}}^i)), \quad (1)$$

$$S_O^i = \text{Cos}(M(\text{comm}_{\text{old}}^i), M(\text{code}_{\text{old}}^i)), \quad (2)$$

$$S_N^i = \text{Cos}(M(\text{comm}_{\text{new}}^i), M(\text{code}_{\text{new}}^i)), \quad (3)$$

$$S_1^i = \text{Average}(S_C^i, S_O^i, S_N^i), \quad (4)$$

where comm^i represents the i -th comment, code^i represents the i -th snippet of code, M represents the pre-trained model (i.e., UniXcoder), and Cos represents cosine similarity. We compute the similarity S_C^i between comments, the similarity S_O^i between the old comments and old code, and the similarity S_N^i between the new comments and new code. The average of these three similarities is used as the score S_1 for Type I noise.

Noise existing between comment changes and code changes (Type II). This type of noise is unique to the comment update task and is mainly caused by inconsistencies between comment changes and code changes. We use the word-level *diff* information as the basic information for calculating the semantics, and use it to evaluate the data quality score. We follow the approach proposed by [8] to get the basic *diff* information between comments and code. For the semantics of *diff* information, we concatenate the changed words in the *diff* into a string and use UniXcoder to calculate the semantics embeddings. We utilize the following formula to calculate the score for the semantics of *diff* information.

$$S_2^i = S_D^i = \text{Cos}(M(d_{\text{comm}}^i), M(d_{\text{code}}^i)), \quad (5)$$

where d_{comm}^i denotes the *diff* information among comments in the i -th data sample, while d_{code}^i denotes the *diff* information among source code in the i -th data sample.

Final score. We consider both types of noise mentioned above and calculate a unified score to assess the data quality of comment updating. Specifically, we calculate two scores and determine the maximum score among them as the final score for each data sample. The formal representation is as follows:

$$S^i = \text{Max}(S_1^i, S_2^i), \quad (6)$$

where the S^i represents the final score of the i -th data sample. We use Max to represent that the semantics of the *diff* only needs to be combined with the maximum value of the semantic similarity between comments and code. This is actually intuitive, as once the semantics of changes are similar, or if there is a high similarity between the comments or code, then it is highly likely to be a data sample that fits the update scenario.

3.1.2 Static filter

In our semantic-aware approach, we quantify the internal correlations of data by evaluating the score of each data sample. Building upon this, we adopt a specific filtering mechanism to ensure that only data samples with higher scores are retained, while those with lower scores are excluded. To be specific, we choose the lower quartile (Q_1) as a reference point and filter out data below this threshold by a certain margin. The formal computation is expressed as follows:

$$X_{\text{low}} = Q_1 - k \times \text{IQR}, \quad (7)$$

where X_{low} denotes the final threshold and IQR represents the interquartile range (i.e., the gap between the upper quartile and the lower quartile). The parameter k ($k > 0$) acts as a coefficient that controls the strictness of data filtering. The larger the value of k , the less data is filtered out.

3.2 Dynamic strategy

In our dynamic strategy, we design a loss-aware technique to evaluate the quality of data.

3.2.1 Dynamic training loss

In contrast to the static strategy, the dynamic strategy integrates loss-aware dynamic evaluations into the data filtering process. This means initiating the process by training a base model and utilizing the loss of individual samples as the primary evaluation metric. Within the domain of deep learning, the loss function serves as a pivotal metric of the model's alignment with the dataset. Hence, a model effectively capturing the underlying knowledge within the data typically yields a lower loss, while ineffective learning manifests in higher loss values.

The advantage of the dynamic strategy lies in its ability to leverage the characteristics of the model, thereby better adapting to the complexity and diversity of the data. We employ the common cross-entropy loss function in deep learning as the approach for calculating loss during model training. To better evaluate the performance of the base model, we choose the top three epochs where the model performs best on the validation set to collect the loss values of data samples. We do this for two main reasons. First, the validation set is typically used to assess the model's generalization ability on unseen data. Therefore, we select epochs based on their performance on the validation set to ensure that the collected data samples are representative and possess good generalization ability. Second, models may encounter issues such as overfitting or underfitting during training, resulting in a decrease in their generalization ability. Thus, we choose results from epochs where the model has converged and demonstrated stable performance.

3.2.2 Dynamic filter

In the loss-aware approach, we take into account both the magnitude and fluctuation of the training loss. This dual consideration is important because samples with consistently high loss and significant variance are likely indicative of cases where the model has failed to learn effectively. For loss magnitude, we simply identify samples whose loss exceeds the upper quartile within an epoch as candidate filtered samples, denoted as C_1 . For loss fluctuation, we compute the variance of each sample's loss across multiple epochs. Samples with loss variance above the upper quartile are similarly marked as candidates, denoted as C_2 . To fulfill our data cleaning objective, we remove samples that appear in both C_1 and C_2 , as these are most likely to hinder effective model training.

3.3 Model ensemble

To integrate the distinct advantages of data cleaned using different strategies, we train two models: M_s , based on data filtered by a static strategy, and M_d , based on data filtered by a dynamic strategy. We then select the final output by comparing their weighted performance scores, formally defined as

$$I = \begin{cases} M_s(x), & \text{if } \alpha \cdot S(M_s) \geq \beta \cdot S(M_d), \\ M_d(x), & \text{otherwise,} \end{cases} \quad (8)$$

where I denotes the selected model output for a given input x ; α and β are weights reflecting each model's performance on a held-out dataset; and $S(\cdot)$ is a scoring function consistent with that used in the static strategy. This approach ensures that the model producing more reliable results under the evaluation metric is chosen at inference time.

4 Experimental setup

In this section, we first introduce the research questions (RQs). Then, we describe the two datasets for the comment updating task and the metrics for evaluating the effectiveness of the updated comments. Finally, we present baselines and experiment settings.

4.1 Research questions

To evaluate the effectiveness and the efficiency of CupCleaner, we aim to answer the following questions.

RQ1: whether CupCleaner can filter out true noise data? To evaluate CupCleaner's ability to remove noise data, we first record the changes before and after the cleaning process. Then we provide examples that CupCleaner identifies as noisy. Finally, we conduct human evaluation on the noisy and high-quality data identified by CupCleaner to explore whether the filtered data are truly noise.

RQ2: how effective does CupCleaner clean the training data? To evaluate the effectiveness of our data cleaning approach, CupCleaner, we conduct experiments on two widely adopted general-purpose code models: PLBART and CodeT5. We apply CupCleaner to clean their training and validation sets, while keeping the test set unchanged across all approaches to ensure fair comparison. This setup ensures that only our approach benefits from data cleaning, while baseline models are trained and validated on noisy data. We compare the performance of models trained on CupCleaner-cleaned data against multiple baselines. These include both deep learning-based approaches (e.g., CoditT5) and non-deep learning filtering heuristics that rely on shallow matching of comment and code edits. Through these comparisons, we aim to assess whether CupCleaner can more effectively identify and remove noisy samples, thereby enhancing model performance on the comment updating task.

RQ3: how effective is the model ensemble? Model ensemble refers to the technique of combining predictions from multiple models to enhance overall performance. To validate the effectiveness of the model ensemble, we first separately evaluate the static and dynamic cleaning strategies. Then, we evaluate the effectiveness of the model ensemble approach by comparing the performance of individual models and the ensembled model on the test set, confirming the impact of individual strategies on the final model performance.

RQ4: how efficient is CupCleaner? Calculating the quality score for each data sample involves using pre-trained models and training base models. This also brings a potential risk that the time required to compute the scores for all data samples may be unacceptable. To answer this question, we calculate the time consumption for cleaning the entire dataset, including the time required by the static strategy and the dynamic strategy. We chose the time it takes to train a model for one epoch as a baseline for comparing the time required to clean the data.

Table 1 Dataset statistics. ACL: the Annual Meeting of the Association for Computational Linguistics. AAAI: the AAAI Conference on Artificial Intelligence.

Name	Panthaplackel et al., 2020	Panthaplackel et al., 2021
Year	2020	2021
Published	ACL	AAAI
Language	Java	Java
#Pairs	7239	18522
Train/valid/test	5791/712/736	16494/1878/150

4.2 Datasets

To evaluate the effectiveness of our data cleaning approach, we conduct experiments on two datasets: Panthaplackel et al., 2020 [8] and Panthaplackel et al., 2021 [9]. Table 1 summarizes the key statistics of these datasets. Both datasets consist of paired old and new comment-code examples that potentially exhibit update relationships. The input to the model includes the old comment, old code, and the modified new code, while the expected output is the updated comment.

Specifically, the Panthaplackel et al., 2020 dataset contains approximately 7.2k samples, collected from the commit histories of open-source Java projects on GitHub, focusing exclusively on changes to `@return` comments. In contrast, the Panthaplackel et al., 2021 dataset is originally constructed for comment consistency detection, and includes a broader range of comment types such as `@param`, `@return`, and summary comments. Although not initially designed for comment updating, recent work [23] has repurposed a subset of this dataset by extracting instances where comment modifications occurred, discarding examples with unchanged comments. The resulting dataset comprises approximately 20k usable samples for comment update research.

Both datasets employ simple filtering heuristics to select examples that likely reflect meaningful comment-code updates. For example, the Panthaplackel et al., 2020 dataset excludes purely stylistic changes (e.g., spelling corrections or reformatting), while the Panthaplackel et al., 2021 dataset filters out samples with minimal comment edits or code that cannot be parsed into an AST. However, due to the lack of deeper analysis into the nature of comment updates, both datasets still contain a considerable amount of noise, which motivates the need for more rigorous data cleaning strategies such as the one we propose.

4.3 Metrics

To evaluate the performance of the models, we select five representative evaluation metrics in the field of text generation and updating, including XMatch, BLEU, METEOR, SARI, and GLEU [21–23, 35, 39].

4.3.1 XMatch

The first metric is exact match, i.e., the percentage of generated comments that are identical to the ground truth at the string level.

4.3.2 BLEU

The second metric we use is the BLEU score. BLEU [42] is originally used to evaluate the performance of machine translation. BLEU [42] is commonly used to evaluate the similarity of generated code from a lexical perspective [19, 43]. Specifically, we use BLEU-4 as the BLEU score.

4.3.3 METEOR

Our third metric is METEOR [44]. METEOR utilizes word alignment to align the model-generated results with the reference texts and evaluate the generated results. Compared to BLEU, METEOR does not rely on n-gram matching, but instead places more emphasis on the syntax, word order, and semantics of the generated results.

4.3.4 SARI

Our fourth metric is SARI (system for automatic evaluation of text simplification). SARI [45] is an evaluation metric based on edit distance, originally designed to assess the overall quality of text simplification. SARI compares the generated results with the reference in terms of edit distance, as well as the matching of words, phrases, and sentences at various levels, to provide an overall assessment of the quality of the generated results.

4.3.5 GLEU

Our last metric is GLEU [46] (generalized language evaluation understanding). GLEU is a variant of the BLEU metric originally designed for evaluating grammatical correction systems. By rewarding correct edits and penalizing ungrammatical ones, GLEU is closer to human-level judgment than BLEU and is more suitable for tasks involving edits. GLEU captures fluency and grammatical constraints through the use of n-grams, making it more suitable for evaluating tasks involving edits.

4.4 Baselines

To comprehensively evaluate the effectiveness of our proposed method, CupCleaner, we compare it against two types of baselines: (1) widely used pre-trained code models, and (2) existing data cleaning approaches tailored for code-comment tasks. The former comparison aims to assess whether applying CupCleaner can enhance the performance of general-purpose models (e.g., PLBART, CodeT5) beyond their original results on noisy data. The latter comparison focuses on evaluating whether CupCleaner outperforms prior cleaning strategies in identifying high-quality training samples and improving downstream performance. This dual-perspective evaluation allows us to holistically measure the practical impact and general utility of CupCleaner.

4.4.1 Baseline models

To ensure generalizability, we evaluate CupCleaner using widely adopted and model-agnostic pre-trained code models rather than task-specific architectures. We select three recent and representative models: PLBART [21], CodeT5 [22], and CoditT5 [23].

PLBART. PLBART is a sequence-to-sequence model based on the BART [47] architecture and pre-trained on large-scale programming language data, including Java and Python from GitHub and natural language descriptions from StackOverflow. It uses denoising pre-training objectives such as token masking, deletion, and infilling. PLBART has demonstrated strong performance on multiple code understanding and generation tasks, outperforming prior models such as CodeGPT [48] and CodeBERT [39].

CodeT5. CodeT5 extends the T5 [49] encoder-decoder architecture for programming language tasks. It introduces code-aware pre-training objectives, including masked span prediction, identifier tagging, masked identifier prediction, and bimodal dual generation. These tasks help the model better capture the structural and semantic characteristics of code, leading to state-of-the-art results across various code-related tasks [7, 35, 50].

CoditT5. CoditT5 [23] is built on top of CodeT5 and is tailored specifically for code editing tasks. It employs a denoising pre-training framework that perturbs code snippets based on downstream edit distributions. These perturbations include insertion, deletion, or replacement of tokens. The model learns to recover the original input and predict the editing process. CoditT5 has shown superior performance on annotation update tasks compared to more general models.

4.4.2 Baseline cleaning approaches

To further validate the effectiveness of CupCleaner, we compare it with two recent code-specific data cleaning strategies that aim to improve training data quality in code-comment tasks.

CommClean. CommClean applies a set of heuristic rules to filter out low-quality code-comment pairs, particularly those with clearly non-informative or automatically generated comments. It is designed to improve the training of comment generation models. When applicable, samples with updated comments are split into independent before/after pairs to enable filtering.

OCDData. OCDData focuses on identifying meaningful comment updates. It first uses heuristic rules to discard samples where the change between two versions of a comment is minimal, assuming that substantial changes indicate more reliable update examples. These filtered samples are then used to train a discriminator to detect outdated comments, with unchanged or unrelated examples treated as negative cases.

4.5 Implementation settings

The main experimental environment is an Ubuntu OS with two Intel Xeon CPUs, 188 GB RAM, and a NVIDIA TITAN RTX GPU with 24 GB of memory. To save training time and resources, we set the maximum number of epochs for training to 20 and implement an early stopping strategy. Specifically, we implement a common strategy during the training process, where we terminate the training process if the experimental performance on the validation set does not improve within three consecutive epochs. For each dataset, we use exactly the same model settings across all variants of the dataset. In addition, to make the experimental process more generalizable,

Table 2 Data statistics after cleaning.

Dataset	Panthaplackel et al., 2020			Panthaplackel et al., 2021		
	# Samples	# Old/new comment tokens	Edit distance	# Samples	# Old/new comment tokens	Edit distance
Original	5791	10.84/11.79	5.99	16494	13.94/14.70	6.27
Static	5227 (−9.74%)	11.24/12.23	5.65	14922 (−9.53%)	14.31/15.13	5.94
Dynamic-PLBART	4624 (−20.15%)	10.47/10.17	4.35	12410 (−24.76%)	13.47/13.34	4.50
Dynamic-CodeT5	4764 (−17.73%)	11.22/11.77	5.01	13427 (−18.59%)	14.20/14.27	4.72

we simply concatenate the inputs into a long sequence without performing any special processing. In the static cleaning strategy, to filter out lower-quality data as much as possible, we set the coefficient k to 0.5. In the model ensemble, we utilize the commonly used metric BLEU as the metric for assigning performance weights to each model.

5 Results

In this section, we present the experimental results and answer each research question.

5.1 Noise data identified by CupCleaner (RQ1)

To answer this question, we explore the data filtered out by CupCleaner. First, we present the statistical information of data cleaning, then we analyze some samples of low-quality data that CupCleaner identifies, and finally, we conduct a human evaluation of the identified clean data and noisy data.

Table 2 compares the training data before and after different cleaning strategies across two datasets. The static strategy removes about 9%–10% of samples, but key metrics such as comment length and edit distance remain largely unchanged. This indicates that static rules filter obvious noise while preserving the original data distribution. In contrast, dynamic strategies show greater variation across models. PLBART filters over 20%–24% of samples, while CodeT5 removes around 18%, reflecting differences in model capacity and focus. Since dynamic filtering relies on model confidence, the retained samples tend to have clearer and more meaningful edits, leading to a larger drop in average edit distance compared to static filtering. Therefore, static methods ensure distributional consistency, while dynamic methods favor higher-quality edits, resulting in more pronounced shifts in the dataset.

To demonstrate the discriminating ability of CupCleaner on low-quality data, we provide a data sample analysis. Figure 3 illustrates the low-quality data samples detected by the static strategy and the dynamic strategy applied on CodeT5, respectively. Figure 3(a) shows an example where the old code contains many lines, but its comment is inadequate for describing the content of the code. Most of the code lines are removed in the edit from the old code to obtain the new code, but the comments in the new code are weakly related to the comments in the old code and are unsuitable for describing the new code adequately. Therefore, the static strategy filters out this data as low-quality data. Figure 3(b) showcases an example where the dynamic strategy identifies cases where the code content remains largely unchanged, but the comment undergoes significant changes. Moreover, the new comment is deemed ineffective in describing the new code. Hence, such data may mislead the model regarding the modification of old comments, resulting in considerable loss during training.

To further assess whether CupCleaner inadvertently filters out valuable data or introduces distributional bias, we conducted a human evaluation focused on the authenticity of samples removed by the static cleaning strategy. Three Ph.D. students in computer science volunteered to participate in the evaluation. This analysis helps verify whether the filtered data indeed lacks semantic alignment between code and comment changes.

We first randomly select 50 samples from the high-quality data and 50 samples from the noisy data identified by CupCleaner in the dataset of Panthaplackel et al., 2020. Due to the susceptibility of the dynamic strategy to different base models, we only investigated whether the data filtered out by the static strategy truly contains noisy data. Then, we make sure that the evaluators understand the context and the intention of the comment updating task and ask them the question “Would you accept this data sample into the comment updating dataset?”. This question means that the annotators need to judge whether the comment update process of each data is reasonable and understandable to humans. The label for each data sample ranges from 1 to 5, representing reject, weak reject, weak accept, accept, and strong accept, indicating the willingness to accept the data.

To ensure objectivity, we assign corresponding characteristics to each score for reference.

- Score 1. Comments are mostly unreasonable or non-literal.
- Score 2. Comments are generally reasonable, but the changes in comments are unrelated to changes in the code.

```

// intent
public static intent get install app intent ( final file file , final string authority ) {
    if ( file == null ) return null ;
    intent intent = new intent ( intent . action _ view ) ;
    uri data ; string type = str ;
    if ( build . version . sdk _ int < build . version _ codes . n ) {
        data = uri . from file ( file ) ;
    } else {
        intent . set flags ( intent . flag _ grant _ read _ uri _ permission ) ;
        data = file provider . get uri for file ( utils . get context ( ) ,
            authority , file ) ; }

    intent . set data and type ( data , type ) ;
    return intent . add flags ( intent . flag _ activity _ new _ task ) ;
}
// app ( 8 . 0 )
public static intent get install app intent ( final file file , final string authority ) {
    return get install app intent ( file , authority , false ) ;
}

```

(a) An example of data filtered out by the static strategy.

```

// true if the string is not empty and not null and not whitespace
public static boolean is not blank ( string str ) {
    return ! is blank ( str ) ;
}
// true : / false :
public static boolean is not blank ( string string ) {
    return ! is blank ( string ) ;
}

```

(b) An example of data filtered out by the dynamic strategy.

Figure 3 (Color online) Noise samples detected by the static strategy and the dynamic strategy, respectively.

- Score 3. Comments are reasonable, and the changes in comments are loosely related to changes in the code.
- Score 4. Comments are reasonable, and the changes in comments are mostly related to changes in the code.
- Score 5. Comments are reasonable, and the changes in comments are strongly related to changes in the code.

We gather all volunteer evaluators to collectively review and discuss the scores of each sample. The final score is determined through a voting process involving the participation of three volunteer evaluators. Finally, the average score of data that CupCleaner considers to be retained is 4.3, while the average score of data filtered out by CupCleaner is 2.5. This indicates that the majority of the data filtered out by CupCleaner belongs to noisy data, while the data retained is predominantly in line with the intent of the comment updating task.

Answer to RQ1. CupCleaner can filter out a significant amount of data (i.e., CodeT5’s dynamic strategy filters out approximately 18% of the data). The results of the human evaluation show that the average score of the data filtered out by CupCleaner is only 2.5/5, while the average score of the high-quality data identified by CupCleaner as high-quality is 4.3/5. This indicates that CupCleaner can distinguish noisy data from high-quality data in comment updating datasets.

5.2 Effectiveness of training using cleaned data (RQ2)

Table 3 presents a comprehensive evaluation of our proposed method, CupCleaner, from both the model improvement perspective and the data cleaning effectiveness perspective. This dual-perspective analysis allows us to holistically assess CupCleaner’s practical value. We use the superscript α to indicate that the difference between baseline and CupCleaner is not statistically significant. Importantly, only our method involves cleaning the training and validation data, while all baseline approaches are trained and validated on the original data. In all experiments, we keep the test set unchanged to ensure fair and consistent evaluation across approaches.

From the model performance perspective, we aim to evaluate whether applying CupCleaner to general-purpose models (PLBART, CodeT5) can enhance their performance on noisy real-world data. The results show clear and consistent improvements. On the Panthaplackel et al., 2020 dataset, CupCleaner-PLBART achieves over 32% gain in XMatch (i.e., from 12.23 to 16.17) and over 22% gain in BLEU (i.e., from 39.7 to 48.82), compared to the original

Table 3 Experimental results.

Dataset	Model	XMatch	BLEU	Meteor	SARI	GLEU
Panthaplackel et al., 2020	PLBART	12.23	39.7	35.74	39.87	34.27
	CodeT5	21.60	55.42 ^α	48.76	44.46	49.49 ^α
	CoditT5	17.39	52.65	45.68	41.64	46.39
	CommClean-CodeT5	21.88	55.15 ^α	48.91	45.16	49.30
	OCDData-CodeT5	21.60	55.23 ^α	48.81	46.02 ^α	49.27
	CupCleaner-PLBART	16.17	48.82	43.40	42.11	42.17
	CupCleaner-CodeT5	23.10	55.83 ^α	49.54	46.63 ^α	50.30 ^α
Panthaplackel et al., 2021	PLBART	33.33	59.68	55.70	55.23	55.08
	CodeT5	41.33	67.03	61.45	63.30	60.54
	CoditT5	43.33	64.56	60.75	61.41	59.53
	CommClean-CodeT5	41.33	67.64	61.56	63.03	60.65
	OCDData-CodeT5	40.67	67.28	60.85	63.93	59.77
	CupCleaner-PLBART	38.00	65.06	59.24	57.37	57.50
	CupCleaner-CodeT5	45.33	69.11	64.37	65.82	62.52

PLBART trained on unfiltered data. Similarly, CupCleaner-CodeT5 improves the already strong CodeT5 model by 6.9% in XMatch (i.e., from 21.60 to 23.10) and achieves the highest scores across all metrics. On the Panthaplackel et al., 2021 dataset, CupCleaner-PLBART improves XMatch by 14.0%, while CupCleaner-CodeT5 again outperforms all baselines, including CoditT5, with a 4.0 percentage points gain in XMatch, and over 2.0 points gain in BLEU (i.e., from 67.03 to 69.11). For most metrics, the improvements brought by CupCleaner are statistically significant, and even in cases marked with α , it still achieves consistent numerical gains over all baselines. These results demonstrate that high-quality training data contributes more significantly to final performance than architectural specialization alone: even a general-purpose model trained on cleaned data can outperform a code-editing-specific model like CoditT5.

From the data cleaning effectiveness perspective, we compare CupCleaner with existing heuristic-based filtering approaches, CommClean and OCDData, applied to CodeT5. As shown in the table, both methods provide minimal or no meaningful gains. On the Panthaplackel et al., 2020 dataset, CommClean improves XMatch only marginally (i.e., from 21.60 to 21.88), while BLEU drops slightly (i.e., from 55.42 to 55.15). OCDData performs similarly, with unchanged XMatch and decreased BLEU (55.23). On the Panthaplackel et al., 2021 dataset, CommClean and OCDData improve BLEU by less than 1 point, and in some cases, GLEU and SARI even drop. In contrast, CupCleaner not only achieves consistently higher scores on all metrics but also introduces a more principled and scalable cleaning mechanism: a static filter based on semantic rules and a dynamic filter guided by model loss.

According to our further analysis, CommClean and OCDData each remove only about 1% of the data, and in some cases, as little as 0.26% (e.g., by CommClean on the second dataset). By comparison, CupCleaner filters out over 9% of the data through its static filtering mechanism, and removes over 17% of noisy samples via its dynamic loss-aware strategy. This demonstrates that existing heuristic rules are too conservative to identify noisy training data effectively, especially in the context of comment updates.

In summary, these results confirm that CupCleaner is both an effective data quality enhancer and a robust training-time plug-in for improving code-comment models. Its ability to deliver consistent improvements across multiple datasets, models, and metrics highlights its generality, while its clear outperformance of prior cleaning methods underscores its practical utility in real-world noisy-code scenarios.

Answer to RQ2. Models trained on data cleaned with CupCleaner consistently outperform those trained on unfiltered data, surpassing even specialized models like CoditT5 across all metrics. For example, on the Panthaplackel et al., 2021 dataset, the BLEU score of PLBART improves from 59.68% to 65.06%, and CodeT5 improves from 67.03% to 69.11%. Compared to existing cleaning methods like CommClean and OCDData, which filter less than 1% of data and yield marginal gains, CupCleaner removes over 9% of noisy samples and delivers significant performance improvements. These results confirm CupCleaner’s effectiveness in enhancing data quality and model performance in comment updating tasks.

5.3 Effectiveness of model ensemble (RQ3)

To answer RQ3, we conduct an ablation study to evaluate the individual impact of the static and dynamic cleaning strategies, as well as the effect of their ensemble. Table 4 reports the results on both datasets, highlighting the first three representative evaluation metrics (XMatch, BLEU, and Meteor), with other metrics showing similar trends.

Table 4 Model ensemble.

Dataset Model	Panthaplackel et al., 2020			Panthaplackel et al., 2021			
	XMatch	BLEU	Meteor	XMatch	BLEU	Meteor	
CupCleaner-PLBART	Ensemble	16.17	48.82	43.4	38.00	65.06	59.24
	Static	14.67	45.73	41.12	36.00	62.87	57.46
	Dynamic	14.81	44.45	39.70	38.00	63.33	58.91
CupCleaner-CodeT5	Ensemble	23.10	55.83	49.54	45.33	69.11	64.37
	Static	23.51	55.78	49.71	44.00	67.89	63.28
	Dynamic	22.15	54.46	48.59	44.67	68.98	63.83

Table 5 The time consumption.

	Panthaplackel et al, et al., 2020	Panthaplackel et al, et al., 2021
Training one epoch	1 epoch (4.38 min)	1 epoch (12.15 min)
Static	~0.5 epoch	~0.5 epoch
Dynamic	~ n epochs	~ n epochs

We observe that both static and dynamic strategies lead to performance improvements over models trained on the original unfiltered data. For example, with the CupCleaner-PLBART model on the Panthaplackel et al., 2020 dataset, static and dynamic strategies improve XMatch from 12.23% to 14.67% and 14.81%, respectively. On the Panthaplackel et al., 2021 dataset, static cleaning boosts XMatch from 33.33% to 36.00%, while dynamic cleaning raises it further to 38.00%. Similar trends are observed for CupCleaner-CodeT5: on the 2020 dataset, static cleaning yields slightly better XMatch (23.51%) than dynamic (22.15%); whereas on the 2021 dataset, dynamic cleaning (44.67%) slightly outperforms static (44.00%).

While the relative effectiveness of the two strategies may vary by dataset and model, the ensemble approach generally offers the best overall performance. For example, CupCleaner-CodeT5 on the 2021 dataset achieves the highest XMatch (45.33%), BLEU (69.11%), and Meteor (64.37%) when combining both strategies. CupCleaner-PLBART also benefits from ensembling, with BLEU improving from 63.33% (dynamic only) to 65.06%, despite XMatch remaining the same. These results indicate that the ensemble mechanism effectively integrates complementary strengths from both static and dynamic strategies.

Notably, the ensemble selection mechanism adopts the same criteria as used in the static cleaning strategy to rank candidate outputs. This consistency further reinforces the effectiveness and reliability of the static component.

Answer to RQ3. Both static and dynamic cleaning strategies individually improve model performance across datasets. Their combination through model ensemble further enhances generalization by leveraging the complementary benefits of each strategy. For example, on the Panthaplackel et al., 2021 dataset, CupCleaner-CodeT5 achieves higher scores across all metrics with ensembling (e.g., BLEU improves from 67.89% to 69.11%, and Meteor from 63.28% to 64.37%), demonstrating the effectiveness of the combined approach.

5.4 Time consumption of data cleaning process (RQ4)

To answer RQ4, we examine the time required to clean the dataset and compare it to training one epoch of the model on the raw dataset. An epoch means that the model goes through all the training data once and adjusts its weights based on the loss function during the training process. We select CodeT5 as the baseline model to calculate the time required to train one epoch, as it has exhibited relatively high performance in previous experiments. In one epoch of model training, we calculate the time spent on the training set and the time spent on the validation set after one epoch. During the data cleaning process, we calculate the time spent on static and dynamic strategies. Table 5 shows the basic statistics of time consumption, where training one epoch on the Panthaplackel et al., 2020 dataset takes approximately 4.38 min, while on the Panthaplackel et al., 2021 dataset, it takes around 12.15 min.

For the static cleaning strategy, the total time consumption is approximately equivalent to 0.5 training epochs, around 2.2 min on the Panthaplackel et al., 2020 dataset and 6.7 min on the Panthaplackel et al., 2021 dataset. This is because generating semantic representations for all training samples only involves the encoder, without requiring full forward-backward propagation, effectively halving the computational workload compared to a full epoch. After this encoding step, the filtering process itself is lightweight and incurs negligible additional time.

The dynamic cleaning strategy requires training the model to converge in order to collect loss signals for filtering, which involves n epochs of training. However, this cost is typically amortized or unavoidable in practice, as model training is necessary regardless of whether cleaning is applied. Moreover, the cost is negligible when compared to

pre-training: for instance, CoditT5 was pre-trained for 4 days [23], while fine-tuning on the 2020 dataset took only 35 min. The loss analysis step itself incurs minimal overhead (i.e., about 0.1 min).

In summary, both the static and dynamic strategies are efficient in practice. Static cleaning is lightweight, and dynamic cleaning reuses already-required training steps. Therefore, the modest additional time required by CupCleaner is acceptable and well justified by the significant improvements it brings to data quality and model performance.

Answer to RQ4. The time cost of CupCleaner is low and practical. Static cleaning requires less time than a single training epoch, while dynamic cleaning leverages standard training steps and adds negligible overhead. Overall, the modest time investment is well justified by the performance and generalization improvements enabled by higher-quality training data.

6 Discussion

Threshold selection and the role of parameter k . In our static semantic-aware filtering strategy, we adopt a thresholding mechanism inspired by the box-plot statistical method. Specifically, we filter out data samples whose semantic similarity scores fall below $Q_1 - k \cdot \text{IQR}$, where Q_1 is the lower quartile and IQR is the interquartile range. This design is motivated by the classical use of box plots to identify statistical outliers. Using Q_1 instead of the median or mean offers robustness to skewed distributions, which are common in real-world code-comment similarity scores, and allows us to focus on the “low-quality tail” of the distribution more effectively.

The parameter k controls the aggressiveness of the filtering strategy: larger values of k result in more conservative thresholds (removing fewer samples), while smaller values make the strategy more aggressive. We conducted further experiments to evaluate the sensitivity of CupCleaner to different k values. The results show that CupCleaner consistently improves model performance across a range of k values (e.g., 0.5, 1.0, 1.5), indicating that the method is robust to the specific choice of threshold. For example, on the Panthaplackel et al., 2021 dataset, the original CodeT5 model (without cleaning) achieves an XMatch score of 41.33% and a BLEU score of 67.03%. When applying CupCleaner with $k = 1.0$, the model improves to 42.00% (XMatch) and 67.49% (BLEU). When the threshold is set to a more conservative value with $k = 1.5$, the scores further increase to 43.33% (XMatch) and 68.44% (BLEU). These results demonstrate that CupCleaner consistently enhances model performance under different filtering intensities, confirming the robustness and practical value of our strategy. In practice, $k = 0.5$ offers a balanced trade-off between recall and precision in identifying noisy samples, which is why we adopt it as the default setting.

Generalizability to other programming languages. Our current evaluation focuses on Java datasets, which aligns with prior work in comment updating [8,9] and reflects Java’s wide adoption in open-source projects. Nonetheless, CupCleaner is language-agnostic by design: its semantic filtering relies on pretrained models’ representations of code-comment pairs, and its dynamic strategy is guided by model loss, both of which apply to other programming languages as long as pretrained encoders are available. While we leave multi-language evaluation for future work, we believe our method can generalize well to other languages, such as Python or JavaScript, especially with the growing availability of multilingual code models.

7 Threats to validity

Threats to internal validity may arise from three factors. The first threat relates to the implementation of our experiments. To mitigate this threat, we load pre-trained model parameters using the HuggingFace [51] framework and employ its standard library to construct the training and evaluation framework for the model. Additionally, we set a random seed to ensure the reproducibility of experiments. Finally, the implementation of model evaluation adopted in our study remains consistent with the existing literature. The second threat is that while filtering out noisy data, we might inadvertently exclude normal data as well. This is easily understandable because defining whether a data sample is noise can be challenging. On one hand, programmers have different ideas and styles when coding, and on the other hand, many data samples require an in-depth understanding to determine their appropriateness. To mitigate this threat, we employ various evaluation metrics from the code generation domain to assess whether models trained on cleaned data can effectively enhance performance and generalization. Experimental results indicate that our cleaning approach significantly improves the quality of the training set and enhances multiple model metrics. The third threat relates to human evaluation. In our human evaluation, we present annotators with only one question and ask them to provide ratings based on their responses to this question. A comprehensive assessment of a data sample may require analysis from multiple perspectives. However, the annotators we invite are all Ph.D students with relevant development experience. To ensure their understanding of

the comment updating task, we also engage in discussions during the annotation process to mitigate the potential impact of this threat.

Threat to external validity may come from the choice of the baseline models. In our experiments, we exclusively choose pre-trained models and do not include non-pre-trained models, which could potentially introduce bias into the final experimental results. We assume that pre-trained models possess a broader base of knowledge, making them more capable of understanding downstream tasks compared to training from scratch. Actually, pre-trained models typically exhibit better performance and are also employed in various generative tasks.

8 Conclusion

In this paper, we propose CupCleaner, a data cleaning approach for comment updating datasets. Our data cleaning approach combines a semantic-aware static cleaning strategy with a loss-aware dynamic cleaning strategy. In the experiments to evaluate CupCleaner, we first conduct a human evaluation, and statistical results suggest that the data identified as noise by CupCleaner receive significantly low scores. Then, we compare models trained on data cleaned with CupCleaner to other code pre-training models. Results show that models trained on cleaned data significantly outperform those trained on original data on the same test set. We also compare the experimental performance of different cleaning strategies and model ensembles. Results indicate that while different cleaning strategies exhibit varied performances on different datasets, they all show improvement compared to models trained on raw data, demonstrating the effectiveness of each cleaning strategy. Furthermore, the model ensemble effectively leverages the characteristics of different strategies, further enhancing model performance and generalization. Finally, we investigate the time consumption of data cleaning and find that the time spent on cleaning data is acceptable compared to fine-tuning models.

In future work, we plan to explore other datasets for various software engineering tasks and expand our data cleaning approach.

Acknowledgements This work was supported by National Key Research and Development Program of China (Grant No. 2023YFB4503803) and National Natural Science Foundation of China (Grant Nos. 62232003, 62402482).

References

- Lientz B P, Swanson E B, Tompkins G E. Characteristics of application software maintenance. *Commun ACM*, 1978, 21: 466–471
- Bennett K H, Rajlich V T. Software maintenance and evolution: a roadmap. In: *Proceedings of the Conference on the Future of Software Engineering*, 2000. 73–87
- Tufano R, Pasarella L, Tufano M, et al. Towards automating code review activities. In: *Proceedings of the 43rd International Conference on Software Engineering (ICSE)*, 2021. 163–174
- Bacchelli A, Bird C. Expectations, outcomes, and challenges of modern code review. In: *Proceedings of the 35th International Conference on Software Engineering (ICSE)*, 2013. 712–721
- Just R, Jalali D, Ernst M D. Defects4j: a database of existing faults to enable controlled testing studies for Java programs. In: *Proceedings of the 2014 International Symposium on Software Testing and Analysis*, 2014. 437–440
- Zhu Q, Sun Z, Xiao Y A, et al. A syntax-guided edit decoder for neural program repair. In: *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2021. 341–353
- Liang Q, Sun Z, Zhu Q, et al. Directional diffusion-style code editing pre-training. 2025. [ArXiv:2501.12079](https://arxiv.org/abs/2501.12079)
- Panthaplackel S, Nie P, Gligoric M, et al. Learning to update natural language comments based on code changes. 2020. [ArXiv:2004.12169](https://arxiv.org/abs/2004.12169)
- Panthaplackel S, Li J J, Gligoric M, et al. Deep just-in-time inconsistency detection between comments and source code. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021. 427–435
- Gao Z, Xia X, Lo D, et al. Automating the removal of obsolete todo comments. In: *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2021. 218–229
- Zha D, Bhat Z P, Lai K H, et al. Data-centric artificial intelligence: a survey. 2023. [ArXiv:2303.10158](https://arxiv.org/abs/2303.10158)
- Chen X P, Hu X, Huang Y, et al. Deep learning-based software engineering: progress, challenges, and opportunities. *Sci China Inf Sci*, 2025, 68: 111102
- Wang G, Sun Z, Gong Z, et al. Do advanced language models eliminate the need for prompt engineering in software engineering? 2024. [ArXiv:2411.02093](https://arxiv.org/abs/2411.02093)
- Yin P, Deng B, Chen E, et al. Learning to mine aligned code and natural language pairs from stack overflow. In: *Proceedings of the 15th International Conference on Mining Software Repositories*, 2018. 476–486
- Munaiah N, Kroh S, Cabrey C, et al. Curating GitHub for engineered software projects. *Empir Software Eng*, 2017, 22: 3219–3253
- Vasilescu B, Serebrenik A, Filkov V. A data set for social diversity studies of GitHub teams. In: *Proceedings of the 12th Working Conference on Mining Software Repositories*, 2015. 514–517
- Rokon M O F, Islam R, Darki A, et al. SourceFinder: finding malware source-code from publicly available repositories in GitHub. In: *Proceedings of the 23rd International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2020)*, San Sebastian, 2020. 149–163
- Sun Z, Li L, Liu Y, et al. On the importance of building high-quality training datasets for neural code search. In: *Proceedings of the 44th International Conference on Software Engineering*, 2022. 1609–1620
- Shi L, Mu F, Chen X, et al. 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. 107–119
- Huang Y, Jia N, Shu J, et al. Does your code need comment? *Softw Pract Exp*, 2020, 50: 227–245
- Ahmad W U, Chakraborty S, Ray B, et al. Unified pre-training for program understanding and generation. 2021. [ArXiv:2103.06333](https://arxiv.org/abs/2103.06333)
- Wang Y, Wang W, Joty S, et al. Codet5: identifier-aware unified pre-trained encoder-decoder models for code understanding and generation. 2021. [ArXiv:2109.00859](https://arxiv.org/abs/2109.00859)
- Zhang J, Panthaplackel S, Nie P, et al. Coditt5: pretraining for source code and natural language editing. In: *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, 2022. 1–12

- 24 Guo D, Zhu Q, Yang D, et al. Deepseek-coder: when the large language model meets programming—the rise of code intelligence. 2024. ArXiv:2401.14196
- 25 Liang Q, Zhang Z, Sun Z, et al. Grammar-based code representation: is it a worthy pursuit for LLMs? 2025. ArXiv:2503.05507
- 26 Liang Q, Zhang Z, Liu C, et al. Condor: a code discriminator integrating general semantics with code details. 2024. ArXiv:2412.17429
- 27 Shi E, Wang Y, Du L, et al. On the evaluation of neural code summarization. In: Proceedings of the 44th International Conference on Software Engineering, 2022. 1597–1608
- 28 Xu S, Yao Y, Xu F, et al. Data quality matters: a case study of obsolete comment detection. In: Proceedings of the 45th International Conference on Software Engineering (ICSE), 2023. 781–793
- 29 Jiang J, Xiong Y, Zhang H, et al. Shaping program repair space with existing patches and similar code. In: Proceedings of the 27th ACM SIGSOFT International Symposium on Software Testing and Analysis, 2018. 298–309
- 30 Long F, Rinard M. Automatic patch generation by learning correct code. In: Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, 2016. 298–312
- 31 Wen M, Chen J, Wu R, et al. Context-aware patch generation for better automated program repair. In: Proceedings of the 40th International Conference on Software Engineering, 2018. 1–11
- 32 Xiong Y, Wang B, Fu G, et al. Learning to synthesize. In: Proceedings of the 4th International Workshop on Genetic Improvement Workshop, 2018. 37–44
- 33 Xiong Y, Wang J, Yan R, et al. Precise condition synthesis for program repair. In: Proceedings of the 39th International Conference on Software Engineering (ICSE). 2017. 416–426
- 34 Sadowski C, Söderberg E, Church L, et al. Modern code review: a case study at Google. In: Proceedings of the 40th International Conference on Software Engineering: Software Engineering in Practice, 2018. 181–190
- 35 Li Z, Lu S, Guo D, et al. Codereviewer: pre-training for automating code review activities. 2022. ArXiv:2203.09095
- 36 Liu Z, Xia X, Yan M, et al. Automating just-in-time comment updating. In: Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering, 2020. 585–597
- 37 Devlin J, Chang M W, Lee K, et al. Bert: pre-training of deep bidirectional transformers for language understanding. 2018. ArXiv:1810.04805
- 38 Liu Y, Ott M, Goyal N, et al. Roberta: a robustly optimized bert pretraining approach. 2019. ArXiv:1907.11692
- 39 Feng Z, Guo D, Tang D, et al. Codebert: a pre-trained model for programming and natural languages. 2020. ArXiv:2002.08155
- 40 Guo D, Ren S, Lu S, et al. Graphcodebert: pre-training code representations with data flow. 2020. ArXiv:2009.08366
- 41 Guo D, Lu S, Duan N, et al. Unixcoder: unified cross-modal pre-training for code representation. 2022. ArXiv:2203.03850
- 42 Papineni K, Roukos S, Ward T, et al. Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, 2002. 311–318
- 43 Liang Q, Sun Z, Zhu Q, et al. Lyra: a benchmark for turducken-style code generation. 2021. ArXiv:2108.12144
- 44 Banerjee S, Lavie A. 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, 2005. 65–72
- 45 Xu W, Napoles C, Pavlick E, et al. Optimizing statistical machine translation for text simplification. *TACL*, 2016, 4: 401–415
- 46 Napoles C, Sakaguchi K, Post M, et al. Ground truth for grammatical error correction metrics. In: Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing, 2015. 588–593
- 47 Lewis M, Liu Y, Goyal N, et al. Bart: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. 2019. ArXiv:1910.13461
- 48 Lu S, Guo D, Ren S, et al. Codexglue: a machine learning benchmark dataset for code understanding and generation. 2021. ArXiv:2102.04664
- 49 Raffel C, Shazeer N, Roberts A, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. 2020. ArXiv:1910.10683
- 50 Liang Q, Sun Z, Zhao Y, et al. Bipartite-grammar aware pretraining for XML-SQL code updating. *ACM Trans Softw Eng Methodol*, 2025, doi: 10.1145/3731752
- 51 Huggingface. Huggingface transformers. 2023. <https://huggingface.co/>