

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

5-2023

Generation-based code review automation: How far are we?

Xin ZHOU

Kisub KIM

Bowen XU

DongGyun HAN

Junda HE

See next page for additional authors

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



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

Citation

ZHOU, Xin; KIM, Kisub; XU, Bowen; HAN, DongGyun; HE, Junda; and LO, David. Generation-based code review automation: How far are we?. (2023). *Proceedings of the 31st IEEE/ACM International Conference on Program Comprehension, Melbourne, Australia, 2023 May 15-16*. 215-226.

Available at: https://ink.library.smu.edu.sg/sis_research/8567

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

Author

Xin ZHOU, Kisub KIM, Bowen XU, DongGyun HAN, Junda HE, and David LO

Generation-based Code Review Automation: How Far Are We?

Xin Zhou
Singapore Management University
Singapore
xinzhou.2020@phdcs.smu.edu.sg

Kisub Kim*
Singapore Management University
Singapore
kisubkim@smu.edu.sg

Bowen Xu
Singapore Management University
Singapore
bowenxu@smu.edu.sg

DongGyun Han
Royal Holloway, University of London
UK
donggyun.han@rhul.ac.uk

Junda He
Singapore Management University
Singapore
jundahe@smu.edu.sg

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

Abstract—Code review is an effective software quality assurance activity; however, it is labor-intensive and time-consuming. Thus, a number of generation-based automatic code review (ACR) approaches have been proposed recently, which leverage deep learning techniques to automate various activities in the code review process (e.g., code revision generation and review comment generation).

We find the previous works carry three main limitations. First, the ACR approaches have been shown to be beneficial in each work, but those methods are not comprehensively compared with each other to show their superiority over their peer ACR approaches. Second, general-purpose pre-trained models such as CodeT5 are proven to be effective in a wide range of Software Engineering (SE) tasks. However, no prior work has investigated the effectiveness of these models in ACR tasks yet. Third, prior works heavily rely on the Exact Match (EM) metric which only focuses on the perfect predictions and ignores the positive progress made by incomplete answers. To fill such a research gap, we conduct a comprehensive study by comparing the effectiveness of recent ACR tools as well as the general-purpose pre-trained models. The results show that a general-purpose pre-trained model CodeT5 can outperform other models in most cases. Specifically, CodeT5 outperforms the prior state-of-the-art by 13.4%–38.9% in two code revision generation tasks. In addition, we introduce a new metric namely Edit Progress (EP) to quantify the partial progress made by ACR tools. The results show that the rankings of models for each task could be changed according to whether EM or EP is being utilized. Lastly, we derive several insightful lessons from the experimental results and reveal future research directions for generation-based code review automation.

I. INTRODUCTION

Modern software development involves numerous software quality assurance activities, such as defect management [1], [2], [3], testing [4], [5], and code review [6], [7], to ensure the quality of software. Code review specifically requires reviewers to assess whether the source code written by authors satisfies both functional (e.g., compilation and testing) and non-functional (e.g., code readability) requirements. Many studies [8], [9], [10], [11], [12], [13], [14] have shown the

outstanding effectiveness of the code review process in removing defects and improving maintainability and readability. In addition, code review also helps to share programming knowledge among developers [7], [13].

The benefits of code reviews have been broadly recognized and it is widely adopted in both proprietary and open-source software projects. However, the benefits come with extensive manual costs in reviewing. In the modern software development process, a large number of code changes are required to be reviewed per month. For example, Microsoft Bing and Linux projects require 3,000 and 500 reviews, respectively, per month [8], [15]. In addition, studies [16], [17], [18], [19] show that program comprehension (one of the major work for code reviewers) takes up as much as half of a developer’s time, especially when the code quality is low [19]. To reduce the developers’ burden, researchers have created techniques [20], [21], [22] to automate code review activities by leveraging deep learning algorithms. Particularly, they focus on revising the submitted code to address the possible flaws in the code (i.e., Code Revision Before Review) [20], [21], [22], writing review comments based on the submitted code (i.e., Review Comment Generation) [22], and revising the submitted code based on the comments written by reviewers (i.e., Code Revision After Review) [20], [22].

Recent approaches [20], [21], [22] targeting the above code review tasks push code review automation to new heights. Tufano et al. [20] proposed an approach namely Trans-Review based on Transformer [23] to complete the code revision generation tasks. They first abstracted the source code by adopting the src2abs tool [24] to reduce the vocabulary size. Then they trained a sequence-to-sequence Transformer on the abstracted code. Thongtanunam et al. [21] proposed AutoTransform, which leverages the Byte-Pair Encoding (BPE) [25] tokenizer and a sequence-to-sequence Transformer to complete the code revision generation task. They highlighted the problem that the revised version of the code might have new identifiers/literals (i.e., new tokens), and it is challenging to generate those new tokens. The BPE approach could help to mitigate such

*Corresponding author. Email: kisubkim@smu.edu.sg

```

Submitted Code
public void addQueues(Queue... queues) {
    this.addQueueNames(collectQueueNames(queues));
}

Ground Truth Code
public final void addQueues(Queue... queues) {
    addQueueNames(collectQueueNames(queues));
}

Prediction of CodeT5
public void addQueues(Queue... queues) {
    addQueueNames(collectQueueNames(queues));
}

```

Fig. 1. A model prediction with a positive contribution though not perfect. Ground truth code that is accepted by reviewers removes the “this” keyword and adds the “final” keyword (as highlighted). The model prediction of CodeT5 only removes the “this” keyword but still improves the initial submitted code.

a problem. Tufano et al. [22] leveraged the Text-To-Text Transfer Transformer (T5) architecture [26] to develop a pre-trained model for code review (i.e., T5-Review) based on 1.4 million code snippets and texts from Stack Overflow¹ and CodeSearchNet [27]. The results showed that the pre-trained T5-Review led to a significant improvement over the non-pre-trained Transformer models in three generation-based code review tasks.

Although ACR methods have been proven to be beneficial, these methods are not compared against each other which may confuse the practitioners about which one is suitable for a specific task. In addition, general-purpose pre-trained models for code like CodeBERT [28], GraphCodeBERT [29], and CodeT5 [30] have been proven to be effective in a wide range of downstream tasks [31], [32] in the field of software engineering such as code search and code summarization. However, the performance of these models on ACR remains unknown.

Existing ACR works heavily rely on a strict evaluation metric called Exact Match (EM): a revised code is considered as correct only if it is identical to the ground truth code. Although Exact Match is useful, it is very strict and only focuses on the number of perfect predictions. Prior studies [33], [34], [35] indicated that the developers did not blindly trust automated tools and their recommendations. For instance, Winter et al. [33] surveyed 386 software developers about their attitudes toward automatic program repairs (APR) tools which are similar to ACR tools. 47.0% of responses highlighted that human decisions are needed when applying recommendations from automated tools and 23.0% of responses do not trust automated tools at all. In addition, the perfect prediction ratios of ACR tools are still low (i.e., 1.2%–23.2%). Thus, we believe that the current ACR tools are targeting to provide high-quality drafts for code reviewers rather than fully replacing them. In this sense, instead of only focusing on the number of perfect predictions (EM), we need to care more about the improvements made in the predictions compared to the initial submitted code on the average level. Figure 1 shows an example that CodeT5 generates a better draft code based

on the submitted code by removing “this.”. Though the generated code cannot meet all requirements of reviewers, it still makes positive progress, which should be taken into account when evaluating the effectiveness of a model. To quantify the progress made in predictions, we introduce a new evaluation metric namely Edit Progress (EP), which measures the reduction ratio of edits needed from the submitted code to the ground truth.

To fill the gaps, we conduct a large-scale comprehensive experiment by comparing the effectiveness of recent ACR tools as well as the general-purpose pre-trained models on a unified benchmark. Our goal is to figure out how far we have reached in generation-based code review tasks and highlight the challenges and potential research directions for better code review automation.

In summary, our study covers three state-of-the-art ACR tools (i.e., Trans-Review [20], AutoTransform [21], and T5-Review [22]) and three general-purpose pre-trained models (i.e., CodeBERT [28], GraphCodeBERT [29], and CodeT5 [30]) on three code review downstream tasks: Code Revision Before Review (i.e., revising the submitted code to address the possible flaws), Code Revision After Review (i.e., revising the submitted code based on the comments written by reviewers), and Review Comment Generation (i.e., writing review comments pointing out problems in the submitted code).

Our experimental results reveal the following key findings:

- 1) *A pre-trained encoder-decoder model, CodeT5, achieves state-of-the-art performance and outperforms all existing ACR techniques in most cases.* We carry out a comprehensive evaluation and discover that 1) T5-Review is the best model among existing ACR approaches in terms of the Exact Match metric; 2) CodeT5 can outperform all existing ACR approaches by a large margin for most cases.
- 2) *Stack Overflow data is helpful to the Review Comment Generation task.* As more pre-training corpus written in natural language are provided in Stack Overflow data, T5-Review which is pre-trained on Stack Overflow data has significantly outperformed all other methods by a large margin.
- 3) *Partial progress also matters.* We observe that a model that can generate the most perfect predictions may not be able to generate the most close-to-perfect predictions, indicating that there may exist much noise on the failure cases of the approaches that can generate many perfect predictions. In addition, the rankings of approaches in EM and EP can be different. This indicates that we may miss tools that can contribute the most positive progress if it ranks not very high in EM. Thus, measuring partial progress is also important for a comprehensive understanding of the ACR tools’ effectiveness.

Overall, our empirical investigation sheds light on the strength and weaknesses of each ACR approach as well as the future research directions for generation-based code review automation. We summarize our contributions as follows:

¹<https://stackoverflow.com/>

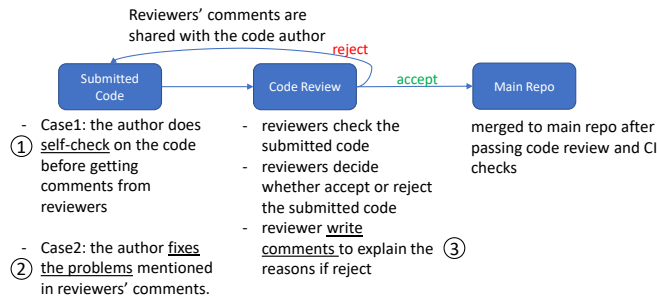


Fig. 2. A brief introduction about the code review processes. The ①, ②, and ③ refer to the three code review activities that generation-based automatic code review tasks automate.

- 1) We conduct a comprehensive investigation on the effectiveness of existing ACR models as well as pre-trained models for code in three downstream code review tasks.
- 2) We introduce a novel evaluation metric to quantify the progress made by ACR tools in generating revised code.
- 3) We derive several insightful lessons from the experimental results and reveal the future direction of the technical design for code review tasks.

II. BACKGROUND

In this section, we introduce the code review processes and formulate the three automatic code review tasks. Then we briefly summarize the automatic code review tools and general-purpose pre-trained code models under investigation.

A. Code Review Processes

Figure 2 briefly shows how code reviews are performed in software development. Code authors first check the correctness of the code by themselves and then submit the code for review. Code reviewers check the submitted code and decide whether to accept/reject to merge the code into the main repository. If the submitted code is rejected, reviewers write comments explaining what and how to improve the code. Then the code authors address those comments and re-submit for review again.

B. Generation-based Code Review Tasks

In this subsection, we describe the generation-based code review downstream tasks. The considered tasks in our study are all proposed to automate the subset of activities in code review processes. We follow the prior studies [20], [21], [22] to formulate each of those tasks.

Code Revision Before Review (CRB) mainly aims to help the code authors who make the code changes to improve and commit them to the repositories. It could help the code authors to address some simple flaws and thus improve the quality of the code *before submitting it for review*. It automates the code review activity ① in Figure 2. At the same time, it also reduces the burden of reviewers on pointing out simple and repeated flaws. This task is formulated as a sequence-to-sequence task: $f(\text{Code}) \rightarrow \text{RevisedCode}$ where f is the DL model and

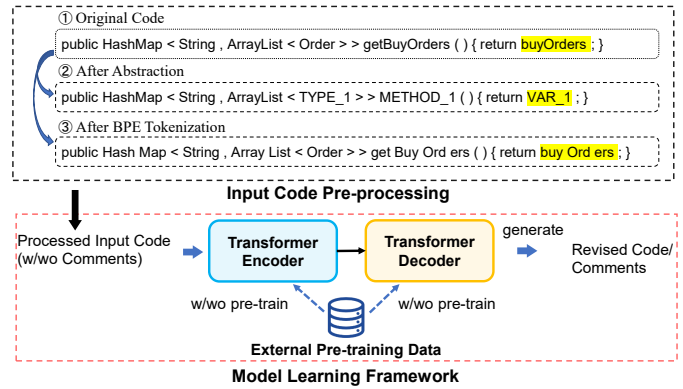


Fig. 3. A Unified Model/Task Framework. All studied models can be fitted in this unified framework by choosing the pre-processing strategy and Transformer models with/without pre-training. All ACR tasks can be fitted into this framework by choosing the corresponding input-output pairs. Note that “w” refers to “with” and “wo” refers to “without”. The parts highlighted in yellow are examples of different pre-processing strategies.

“*Revised Code*” is the ground truth (i.e., the revised code that is accepted by code reviewers).

Code Revision After Review (CRA) aims at supporting code authors by addressing review comments which automate the code review activity ② in Figure 2). The techniques belonging to this task can boost the speed of the revision process as they generate the initial version. This task is formulated as a sequence-to-sequence task: $f(\text{Code}, \text{Comment}) \rightarrow \text{RevisedCode}$ where f is the DL model and “*Revised Code*” is the ground truth code accepted by reviewers. For this task, there have two inputs (i.e., the submitted code and the comments) and a single output (i.e., the revised code).

Review Comment Generation (RCG) is designed for the reviewers. The models in this task draft the initial review comments such that reviewers can save time by revising the written draft, which automates the code review activity ③ in Figure 2. Review Comment Generation is also formulated as a sequence-to-sequence task: $f(\text{Code}) \rightarrow \text{Comment}$ where f is the DL model.

C. Automatic Code Review Tools

All considered tools are built based on a popular Deep Learning model architecture called Transformer [23], which mitigates the long-distance dependency issue of CNN-based models [36] and the long-time computation issue of RNN-based models [23]. A Transformer model typically has two parts: an encoder to understand the semantics of the input data and a decoder to generate the final results (e.g., description, translation, revision, etc.) All studied models adopt the Transformer model (both the encoder and decoder). We highlight the novel designs of each studied model in the following.

Trans-Review [20] are proposed to complete the Code Revision Before Review task (CRB) and the Code Revision After Review (CRA). A large vocabulary size will hinder the learning process and lead to inaccurate generation [37]. To reduce the vocabulary size, Trans-Review abstracted the source

code by adopting the src2abs [24], which replaces actual identifiers/literals in code with reusable IDs. Specifically, for both the submitted and revised versions of code, it replaces identifiers (e.g., types and names of source code elements) and literals (e.g., string values) with an ID that is allowed to be reused across different code snippets. For instance, as shown in Figure 3, the *first* variable appearing in each code will always be replaced with an ID of “1” and a string “VAR” to indicate it is a variable (i.e., VAR_1).

AutoTransform [21] is only designed to complete the Code Revision Before Review task (CRB). Unlike Trans-Review, it applies a Byte-Pair Encoding (BPE) [25] approach to reduce the size of the vocabulary. BPE is a tokenization approach that splits a code token into a sequence of frequently occurring subtokens. For instance, a code token “buyOrders” in Figure 3 will be split into several subtokens “buy”, “Ord”, and “ers”. This strategy brings benefits to efficiency as subtokens are not only occurred more frequently than the regular code tokens, but they also allow to reduce the size of the vocabulary. Empirical evidence showed that BPE could effectively address the large vocabulary size problem [38].

T5-Review [22] leveraged the Text-To-Text Transfer Transformer (T5) architecture [26]. T5 is a pre-training paradigm that corrupts the input data (e.g., randomly masking/dropping/inserting tokens) before feeding into the Transformer encoder and reconstructing the original input data at the Transformer decoder. Those pre-training tasks are called denoising objectives. Tufano et al. pre-trained a code review version of T5 (denoted T5-Review) on 1.4 million code snippets and texts from Stack Overflow and the CodeSearchNet dataset [27]. T5-Review consists of a 6-layer Transformer encoder and a 6-layer Transformer decoder. To reduce the vocabulary size, T5-Review leverages BPE tokenization [25].

D. General-purpose Pre-trained Models for Code

The following techniques are the popular pre-trained models that are trained on source code and texts.

CodeBERT [28] is a SE knowledge-enriched bi-modal pre-trained model, which is capable of modeling both natural languages (NL) and programming languages (PL). CodeBERT leverages a 12-layer Transformer encoder as the model and pre-trained the encoder with the NL-PL data pairs from the CodeSearchNet dataset [27]. It adopts two pre-training objectives jointly: Masked Language Modeling (MLM) [39] and Replaced Token Detection (RTD) [40]. The eventual loss function for CodeBERT is formulated below:

$$\min_{\theta} (\mathcal{L}_{MLM}(\theta) + \mathcal{L}_{RTD}(\theta)) \quad (1)$$

where θ denotes the model parameters of the 12-layer Transformer encoder. CodeBERT has shown great effectiveness in a range of SE downstream tasks such as code search and code summarization [28], [31], [32].

GraphCodeBERT [29] brings the inherent structure of code into consideration during source code modeling. To learn

source code representation, GraphCodeBERT leverages three input components, which are Programming Language, Natural Language, and Data Flow Graph. GraphCodeBERT introduces two structure-aware pre-training tasks (i.e., Edge Prediction and Node Alignment) aside from the MLM prediction task.

CodeT5 [30] is another pre-trained model whose Transformer encoder and decoder are both pre-trained. CodeT5 is capable of establishing complex mappings from sequence to sequence, which is suitable for generation-based tasks. CodeT5 is developed based on the general T5 architecture and is pre-trained not only with the original pre-training tasks (i.e., the denoising objectives) of T5 but also a set of newly proposed pre-training tasks (e.g., Masked Identifier Prediction and Identifier Tagging). Specifically, CodeT5 is enhanced by making it aware of the type information of identifiers.

III. EXPERIMENTAL DESIGN

In this section, we introduce our research questions and describe the corresponding experimental settings.

A. Research Questions

RQ1. Which is the best-performing ACR tool for each targeted code review activity? Recent ACR tools have been demonstrated to be effective in code review tasks. However, a comparative study is missing, which makes it confusing to rank the existing techniques. Addressing this research question may support practitioners in selecting the most suitable technique for each downstream task, and the revealed state-of-the-art that can be considered for the baseline of future research. In this RQ, we conduct a revisiting evaluation of existing ACR tools on the same set of available datasets.

RQ2. Can popular pre-trained models outperform the existing ACR models? General-purpose pre-trained code models have shown significant generalizability and demonstrated promising performance in a broad range of SE tasks [28], [30], [32], [31] with datasets of varying sizes and properties. However, prior ACR tools have not considered these models either as baselines or building blocks of the ACR tools. We believe that discovering the effectiveness of such models on code review tasks would contribute to the research community, and they can also be the baseline for future works. The subject models are listed as follows: (1) CodeBERT [28], GraphCodeBERT [29], and CodeT5 [30]. To answer this research question, we take popular pre-trained models into consideration and compare them with existing ACR models.

RQ3. Comparing the generated code with the human-written code, how far apart are they? The prior ACR works only use a strict evaluation metric (i.e., Exact Match) to analyze the effectiveness of their models: a revised code is considered as a correct generation only if it is identical to the ground truth. Otherwise, no matter if the difference is just a single token or they are different, the generated code is equally considered wrong. Thus, we may miss suggestions that are very close to ground truths that are still partially useful (as shown in the example in Figure 1). It may be interesting to have an overall idea of how close those generated

codes are compared to the ground truths because the current automation of the code review process does not only aim to replace the practitioners but also provide suggestions/drafts for developers to work on [33]. To answer this research question, we introduce a new evaluation metric, Edit Progress, which quantifies the progress achieved by the generated code from the problematic code toward the ground truth. Edit Progress is firstly used in the text editing task to measure the progress [41]. Note that this metric can only be applied in code revision generation tasks because Edit Progress needs an initial input and the revised version to compute whether the revised one improves or degrades the initial one. However, in the Review Comment Generation task, there is no initial comment.

B. Experimental Settings

1) **Datasets**: We include three different datasets used in evaluating the latest ACR tools, namely Trans-Review_{data}, AutoTransform_{data}, and T5-Review_{data}. Tufano et al. [20] collected the Trans-Review_{data} from projects in Gerrit and GitHub. They filtered out noisy comments and comments that are longer than 100 tokens. Besides, they removed the review data where the revised code has new tokens not shown in the initial submitted code. Thongtanunam et al. [21] used the AutoTransform_{data} that is collected from three Gerrit code review repositories: Android², Google³, and Ovirt⁴. Please note that AutoTransform_{data} only has the submitted codes and the revised codes but does not have the corresponding review comments. Tufano et al [22] collected T5-Review_{data} from Java open-source projects from GitHub which have at least 50 pull requests, 10 contributors, 10 stars, and are not forks. They filtered out noisy comments, non-English comments, and duplicate comments.

In this paper, for the Code Revision Before Review (CRB), we employ Trans-Review_{data}, AutoTransform_{data}, and T5-Review_{data}. For the Code Revision After Review (CRA) and Review Comment Generation (RCG), we leverage Trans-Review_{data} and T5-Review_{data}. We do not use AutoTransform_{data} because it does not contain the reviewers’ comments given to submitted codes. The statistics of the used datasets are summarized in Table I.

2) **Evaluation Metrics**: Following previous literature [20], [21], [22], we report the performance of each investigated model using the Exact Match metric (RQ1 and RQ2). We additionally report Edit Progress to measure how close generated codes are compared to the ground truths (RQ3).

Exact Match (EM) is a strict metric. Specifically, for each input-output pair, if the tokens of the model’s prediction (generated by the model given the input) exactly match the tokens of the ground truth (output), EM = 1; otherwise, EM = 0. With just a single token difference between the generated code and the ground truth, the EM score of the generated code is 0.

TABLE I
STATISTICS OF STUDIED DATASETS.

Dataset Statistics	#Train	#Valid	#Test	#New Tokens
Trans-Review	13,756	1,719	1,719	0%
AutoTransform	118,039	14,750	14,750	82.9%
T5-Review	134,239	16,780	16,780	89.6%

*Note that the last column shows the ratio of instances that the ground truth revised code contains new tokens that are not shown in the initial submitted code.

Edit Progress (EP) is a novel metric to measure the progress made by the generated code on the path from the erroneous code to the clean code. Recently, Dibia et al. [42] highlighted the significance of measuring the effort to correct the generated code. They suggested that an edit distance-based metric [43] could serve as an important proxy for estimating the effort required to edit and fix generated code. As an edit distance-based metric, the EP metric may offer a better estimation of the user effort needed to edit code generations in ACR tasks than EM. In automatic code review tasks, each data instance in the test set contains an initial submitted code \tilde{C} and a ground truth code \bar{C} that is revised by developers and accepted by code reviewers. Given the initial submitted code \tilde{C} , each ACR model generates a revised code \hat{C} . The Exact Match (EM) is a strict metric as EM = 1 only when \hat{C} is identical to \bar{C} . While useful, Exact Match also has limitations. It only expects models to be able to fully correct/revise an erroneous code into the correct code. While in some cases, models are still able to make progress by reducing the number of errors in the submitted code, even though the generated code revisions \hat{C} are not identical to the ground truth code revisions \bar{C} . To measure the progress made on the path from the erroneous code to the clean code, we adopt the Edit Progress [41] metric in this work. Specifically, for a data instance (\tilde{C}, \bar{C}) , the Edit Progress (EP) scores of ACR models are computed as the relative edit reduction between the initial code \tilde{C} and the ground truth code \bar{C} by predicting a revision \hat{C} . The metric is formulated in the following equation:

$$Progress = \frac{|D_{\tilde{C} \rightarrow \bar{C}}| - |D_{\hat{C} \rightarrow \bar{C}}|}{|D_{\tilde{C} \rightarrow \bar{C}}|} \quad (2)$$

where D is the edit distance [44] (i.e., the minimum number of operations required to transform one sequence into the other). $|D_{A \rightarrow B}|$ refers to the edit distance between a string A to a string B. An example of how to calculate the EP score is shown in Figure 4. In the example, to turn the submitted code \tilde{C} into the corresponding ground truth code \bar{C} , there are four token-level edit actions: 1) replace “hypot” with “cosh”; 2) delete “;”; 3) delete “double”; 4) delete “y”. If a model generates a prediction like the first prediction in Figure 4, there are three deletions left. Thus, the prediction makes 25% progress on the way from the submitted code to the ground truth revised code. Similarly, the second prediction in Figure 4 makes 75%

²<https://android-review.googlesource.com/>.

³<https://gerrit-review.googlesource.com/>

⁴<https://gerrit.ovirt.org/>




Fig. 4. An Example of the Computation of Edit Progress Scores

progress (reduce 75% edits needed from the submitted code to the ground truth code). A perfect prediction \hat{C} that can fully correct all errors in the submitted code \tilde{C} would achieve a 100% Edit Progress score. A model prediction \hat{C} can have negative progress if it generates more errors than correct edits. We calculate and report the average EP on all test data instances for each model.

Previous studies [43], [45], [42] also consider the edit distance-based similarity metric namely Normalized Edit Distance to measure the similarity between the generated code and ground truth. This metric is the character-level edit distance and can be represented as: $EditDistance(x, y) = \frac{LevenshteinDistance(x, y)}{\max(len(x), len(y))}$ where x and y are strings of the generated code and ground truth, respectively; $len(\cdot)$ is the number of characters in the string; and $LevenshteinDistance(x, y)$ is the number of single-character edits needed to transform one string (x) to another (y). The proposed EP is a novel metric compared to the existing “Normalized Edit Distance” metric:

- 1) EP uses token-level edit distance while Normalized Edit Distance only takes into account single-character edits and may not consider the semantics of code tokens. Furthermore, single-character edits may not be suitable for pre-trained models such as CodeT5, which build their vocabularies on the token or sub-token levels and cannot directly perform single-character edits [46], [47], [48]. EP uses token-level edits, which preserve the meaning of tokens and are feasible for pre-trained models.
- 2) EP compares the generated code with the input code to measure the generation quality of an ACR tool. In ACR tasks, the tools generate revised code, given the input. While the Normalized Edit Distance measures the similarity between the generated code and the ground truth, EP considers not only the similarity but also the improvement of the generated code compared to the input code. This makes EP more suitable than Normalized Edit Distance in situations where the input code is already very similar to the ground truth and needs to be improved further through the generated code.

BLEU score is a widely used metric to measure the similarity of two natural language texts. For review comment generation,

the model generates comments written in natural language. We follow the prior work [49], [50] to adopt the commonly-used metric (i.e., the BLEU score [51]) to measure the quality of generated texts. BLEU scores measure how well a generated text matches the ground truth text by counting the percentage of overlaps.

Please note that BLEU scores are different from Edit Progress scores: BLEU scores measure the similarity between predictions and ground truths, while Edit Progress scores measure the progress the predictions made toward generating the ground truths.

3) **Implementation Details:** For existing ACR tools, we adopt the original replication packages of Trans-Review⁵, AutoTransform⁶, and T5-Review⁷.

By following the standard setting of fine-tuning pre-trained models [39], [28], [29], [30], we set the maximum accepted input sequence as 512, learning rate as 5e-5, and the optimizer as AdamW [52]. For validation, we select the best-performing model checkpoint on the validation set. The best-performing checkpoint is used for testing. We feed the test sets into each approach and get the predictions. Then, we calculate the evaluation metrics based on the predictions and the ground truths.

Both CodeBERT and GraphCodeBERT only have pre-trained encoders. Thus, when applying to a generation-based task where decoders are needed, we follow the original implementations provided by CodeBERT⁸ and GraphCodeBERT⁹ in generation-based tasks to initialize a non-pre-trained decoder. To make the size of the decoder the same as the size of the encoder, we initialize a 12-layer Transformer decoder for CodeBERT and GraphCodeBERT to complete three generation-based code review tasks. We directly adopt the original CodeT5 since it has the pre-trained encode and decoder.

Note that when generating the predictions, we allow each ACR model/pre-trained model to generate *a single prediction (the Top 1 prediction)* given the input data, which is widely used in prior studies [20], [21], [22], [32].

IV. EXPERIMENTAL RESULTS

This section describes the experimental results and answers the research questions. The experimental results are summarized in Tables II, III, and IV, respectively. Note that tools in the light grey are the existing ACR tools and the others are general-purpose pre-trained models. Regarding the performance metrics, we adopt the EM by following prior studies to investigate the effectiveness of the existing ACR tools as well as pre-trained code models in RQ1 and RQ2. We analyze our novel metric, EP, in RQ3 with its comprehensive rationale.

⁵https://github.com/RosaliaTufano/code_review

⁶<https://github.com/aws-research/AutoTransform-Replication>

⁷https://github.com/RosaliaTufano/code_review_automation

⁸<https://github.com/microsoft/CodeBERT/tree/master/CodeBERT/code2nl>

⁹<https://github.com/microsoft/CodeBERT/tree/master/GraphCodeBERT/translation>

TABLE II
EXPERIMENTAL RESULTS FOR THE CODE REVISION BEFORE REVIEW (CRB) TASK

Approaches	Trans-Review _{data}		AutoTransform _{data}		T5-Review _{data}		Average	
	EM	EP	EM	EP	EM	EP	EM	EP
Trans-Review	3.5%	-1.1%	0.1%	-16.6%	0.1%	-151.2%	1.2%	-56.3%
AutoTransform	6.6%	49.7%	10.2%	29.9%	0.8%	9.7%	5.9%	29.8%
T5-Review	7.4%	-14.9%	13.9%	-71.5%	3.3%	13.8%	8.2%	-24.2%
CodeBERT	8.3%	49.8%	8.6%	-75.3%	1.1%	22.3%	6.0%	-1.1%
GraphCodeBERT	6.7%	50.6%	6.1%	-80.9%	0.9%	22.6%	4.6%	-2.6%
CodeT5	8.8%	41.8%	13.7%	-67.8%	5.4%	25.6%	9.3%	-0.1%

The gray rows refer to the existing ACR tools designed for this task and the numbers in bold indicate the best performers in each dataset and metric.

TABLE III
EXPERIMENTAL RESULTS FOR CODE REVISION AFTER REVIEW (CRA) TASK

Approaches	Trans-Review _{data}		T5-Review _{data}		Average	
	EM	EP	EM	EP	EM	EP
Trans-Review	13.5%	52.5%	0.3%	-120.4%	6.9%	-34.0%
T5-Review	24.4%	65.6%	9.0%	38.4%	16.7%	52.0%
CodeBERT	20.2%	57.8%	11.3%	35.2%	15.8%	46.5%
GraphCodeBERT	19.2%	55.9%	11.8%	35.8%	15.5%	45.9%
CodeT5	30.2%	66.8%	16.1%	41.9%	23.2%	54.4%

The gray rows refer to the existing ACR tools designed for this task and the numbers in bold indicate the best performers in each dataset and metric.

TABLE IV
EXPERIMENTAL RESULTS FOR REVIEW COMMENT GENERATION (RCG) TASK

Approaches	Trans-Review _{data}		T5-Review _{data}		Average	
	EM	BLEU	EM	BLEU	EM	BLEU
T5-Review	2.0%	2.5%	2.0%	2.0%	2.0%	2.3%
CodeBERT	0.5%	1.4%	0.8%	0.5%	0.7%	1.0%
GraphCodeBERT	0.3%	1.6%	0.9%	0.6%	0.6%	1.1%
CodeT5	1.1%	1.4%	1.2%	1.8%	1.1%	1.6%

The gray row refers to the existing ACR tool designed for this task and the numbers in bold indicate the best performers in each dataset and metric.

RQ1: Which is the best-performing ACR tool for each targeted code review activity?

For **Code Revision Before Review** (Table II), T5-Review achieves a performance of 7.4%, 13.9%, and 3.3% in terms of Exact Match (EM) on the three datasets, respectively. Among the existing ACR tools, T5-Review shows the best performance.

For **Code Revision After Review** (Table III), the results show that T5-Review performs the best with 24.4% and 9.0% EM scores in Trans-Review_{data} and T5-Review_{data} respectively while Trans-Review achieves 13.5% and 0.3% for the same datasets.

For **Review Comment Generation** (Table IV), among ACR tools, only T5-Review is designed for this task. Thus, we only investigate its performance which achieves 2.0% and 2.3% for EM and BLEU scores on average, respectively.

Overall, we observe that T5-Review shows the best per-

formance among state-of-the-art ACR tools on the studied downstream tasks. The major difference is that T5-Review is pre-trained on external large-scale data while Trans-Review and AutoTransform are not pre-trained. The results indicate that such pre-training is an effective solution to improve performance in most ACR tasks. We present further in-depth investigation on sub-findings (e.g., why Trans-Review struggles) in Section V.

Answer to RQ1: Among the existing ACR methods, T5-Review shows the best performance in the targeted downstream tasks, outperforming the second-best ACR tool by a large margin. It implies that practitioners should consider using T5-Review for ACR tasks.

RQ2: Can popular pre-trained models outperform the existing ACR models?

As we observed that pre-training may bring significant improvement in the case of T5-Review, we investigate further pre-trained code models. Although pre-training a model is usually computationally expensive [26], [28], [30], several general-purpose pre-trained code models (i.e., CodeBERT, GraphCodeBERT, and CodeT5) are publicly released. In this RQ, we study the effectiveness of these pre-trained models under ACR tasks by comparing them against state-of-the-art ACR tools.

We also present the results of the general-purpose pre-trained code models in Table II, III, and IV. As the observation illustrates that the pre-trained models generally outperform non-pre-trained ACR tools, we focus on the comparison between pre-trained code models and T5-Review (i.e., the best existing ACR tool).

For the **Code Revision Before Review** (Table II), on the one hand, CodeBERT and GraphCodeBERT fail to outperform the best ACR approach T5-Review identified in RQ1 for most cases. Specifically, CodeBERT only outperforms T5-Review in Trans-Review_{data} by 12.2% while achieving 61.6% and 200% lower EM scores than T5-Review in the other two datasets. GraphCodeBERT consistently achieves lower EM scores than T5-Review. On the other hand, CodeT5 is able to achieve higher/comparable EM scores of T5-Review. CodeT5 has significantly outperformed T5-Review

in Trans-Review_{data} and T5-Review_{data} datasets by relative improvements of 18.9% ($\frac{8.8-7.4}{7.4}$) and 63.6% ($\frac{5.4-3.3}{3.3}$). However, for the AutoTransform_{data} dataset, CodeT5 achieves a comparable (slightly lower) EM score than T5-Review.

For the **Code Revision After Review** (Table III), CodeT5 still outperforms the best ACR tool, T5-Review, in most cases. Particularly, CodeT5 pushes forward the performances to new heights on EM scores of 30.2% and 16.1%, which are 23.8% ($\frac{30.2-24.4}{24.4}$) and 78.9% ($\frac{16.1-9.0}{9.0}$) over T5-Review improvements. Moreover, CodeBERT and GraphCodeBERT get EM scores of 20.2% and 19.2% for the Trans-Review_{data}, 11.3% and 11.8% for the T5-Review_{data}, which are slightly worse than T5-Review on average.

For the **Review Comment Generation** (Table IV), CodeT5, however, is no longer the best-performing model, as T5-Review shows the best performance. In particular, CodeT5 only achieves 1.1% and 1.6% in EM and BLEU on average, which are crucially lower than T5-Review (i.e., 2.0% and 2.3% in EM and BLEU) on average. Besides, both CodeBERT and GraphCodeBERT struggle with this task obtaining lower scores.

Overall, the results imply that CodeT5 is the most suitable model for Code Revision Before Review (CRB) and Code Revision After Review (CRA), while practitioners can leverage T5-Review for the Review Comment Generation. Further in-depth analysis is provided in Section V.

Answer to RQ2: CodeT5 has impressive performance for two of our target code review tasks (i.e., Code Revision Before Review and Code Revision After Review) by surpassing the best state-of-the-art ACR tool, T5-Review. On the other hand, T5-Review is still the best on a task (i.e., Review Comment Generation). We find that pre-trained models are generally effective for code review tasks.

RQ3: Comparing the generated code with the human-written code, how far apart are they?

In this research question, we adopt the Edit Progress (EP) metric to measure the progress made by each approach on the path from the erroneous code to the golden code. The EP score indicates the percentage of progress made. If EP equals 100%, it means that the generated code has mitigated all the differences between the erroneous code and the ground truth.

For the **Code Revision Before Review** task, as shown in Table II, a model that can generate the most perfect predictions (i.e., EM) may not be able to generate the close-to-perfect predictions (i.e., EP). For the Trans-Review_{data} dataset, the best-performing model in EP is GraphCodeBERT, achieving an EP score of 50.6% while CodeT5 only achieves 41.8%. For the AutoTransform_{data} dataset, AutoTransform shows the best performance with 29.9% in EP. Although CodeT5 performs much better than AutoTransform in EM, CodeT5 achieves much worse performance (i.e., -67.8%) than AutoTransform.

This indicates that CodeT5 generates the most perfect predictions but may generate many noisy failure cases at the same time. While AutoTransform is more conservative: it does not generate many perfect predictions but ensures that most of its revisions bring positive contributions. For T5-Review_{data} dataset, the best model in EP is CodeT5 with a 25.6% EP. This indicates that we need to carefully select an approach for the CRB task when having different needs: either the one that can generate the most perfect predictions (i.e. in terms of EM) or the one that can make the most average positive contributions (i.e., in terms of EP).

Compellingly, Table III reports that CodeT5 consistently outperforms all the other approaches for the **Code Revision After Review** task. It obtains EP scores of 66.8% and 41.9% in both datasets while T5-Review shows comparable results (i.e., 65.6% and 38.4%).

Overall, the results suggest that CodeT5 is still the best model in EP for Code Revision After Review, while AutoTransform generates better code for Code Revision Before Review. Appealingly, AutoTransform achieves a 29.8% EP score on average for Code Revision Before Review, while other approaches show negative progress, which indicates that after their revisions, the generated codes are even worse than the originally submitted code by authors on average. We further verify the research value of EP in Section V.

Answer to RQ3: In summary, the EP metric can change the rank of the approaches for different code review tasks. For instance, in the Code Revision Before Review task, the best model in EP (i.e., AutoTransform) is not even in the top 3 in EM. This also implies there may exist much noise on the failure cases of the approaches that can generate many perfect predictions. CodeT5 can be the best candidate for Code Revision After Review according to the EP metric, while AutoTransform interestingly shows good performance in Code Revision Before Review for the partial progressive aspect.

V. DISCUSSION

This section presents the lessons learned from our experimental results and the threats to validity.

A. Lessons Learned

BPE can significantly improve model performances on CRB and CRA tasks if there are new tokens introduced in ground truth codes. Prior ACR tools (as well as pre-trained code models) adopt either BPE or abstraction to reduce the size of vocabulary for the ACR datasets. Trans-Review, the only ACR tool using abstraction, performs poorly (close to zero) on AutoTransform_{data} and T5-Review_{data} that contain new tokens (not shown in the submitted code) in ground truth code. The abstraction technique replaces actual identifiers/literals in the *submitted code* with reusable IDs to form the vocabulary. However, if there are any new identifiers/literals (have not

TABLE V
RATIOS OF INTERROGATIVE SENTENCES AND THE MODEL PERFORMANCE FOR RCG TASK.

	CodeT5	T5-Review	Diff (%)
Pretraining Data Statistics			
Pretraining Data Source	CodeSearchNet	Stack Overflow	-
Interrogative Sentences (%)	33.8%	69.7%	+106.2%
Review Comment Generation (T5-Review_{data})			
Subsets w/ Interrogative Sentences (EM)	0.24	0.84	+250.0%
Subset wo/ Non-interrogative Sentences (EM)	1.60	3.02	+88.8%

appeared in the submitted code) in the *ground truth revised code*, the built vocabulary does not contain the IDs for new identifiers/literals. Thus, it is challenging for a DL model to generate new tokens beyond the vocabulary. While by using Byte-Pair Encoding (BPE), all other ACR tools (e.g., CodeT5, T5-Review, and AutoTransform) show good performance on datasets with new tokens, outperforming Trans-Review by at least 700% in EM.

The pre-trained encoder-decoder framework significantly outperforms the pre-trained encoder framework by 17.0%–102.2%. Firstly, we find that pre-trained models perform better than not pre-trained models in most cases. For instance, pre-trained T5-Review and CodeT5 consistently show better performance than Trans-Review and AutoTransform in terms of Exact Match. Secondly, we find that within pre-trained models, pre-trained encoder-decoder models (e.g., CodeT5) perform significantly better (17.0%–102.2%) than pre-trained encoder-only models (e.g., CodeBERT and GraphCodeBERT) in generation-based code review tasks. The possible reason may be: when applying to generation-based tasks, pre-trained encoder-only models have to initialize a non-pre-trained decoder. This non-pre-trained decoder will lead to a worse performance [53]. Lastly, within pre-trained encoder-decoder models, general-purpose pre-trained model CodeT5 shows better performance than task-specific pre-trained model T5-Review in most cases. In particular, CodeT5 could outperform T5-Review by 13.4% to 38.9% in terms of EM for two code revision tasks.

Stack Overflow data is helpful to the Review Comment Generation task, especially for comments written in interrogative sentences (250.0% improvements). CodeT5 fails to outperform T5-Review for the Review Comment Generation task. To explain why CodeT5 is inferior to T5-Review in this task, we further investigate the pre-training and downstream task datasets: Firstly, we find that many comments in the Review Comment Generation datasets are interrogative sentences (e.g., “why is there a newArrayList?”). Then we investigate the ratios of interrogative sentences in pre-training data by a simple rule-based method: if a comment contains a question word such as “which”, “where”, “what”, “which”, “does”, we

assume this comment as an interrogative sentence. As shown in Table V, CodeT5 is mainly pre-trained on the CodeSearchNet dataset, which contains 33.8% interrogative sentences while T5-Review is also pre-trained on the Stack Overflow data¹⁰ with a higher ratio of interrogative sentences (i.e. 69.7%). Table V also presents the performances on subsets with/without interrogative sentences, using the T5-Review_{data} as an example. We observe that with a higher ratio of interrogative sentences in the Stack Overflow data, T5-Review can outperform CodeT5 by 250.0% on the subset with interrogative sentences. While T5-Review only outperforms CodeT5 by 88.3% on the subset without interrogative sentences. The fewer interrogative sentences (comments with sentiments) in the CodeSearchNet dataset may hinder CodeT5 in generating interrogative sentences (or comments with sentiment words) given the submitted code.

Overall, we find that Stack Overflow data is helpful to the review comment generation task (improving both with and without interrogative sentence subsets) possibly because it provides more pre-training corpus written in natural language. In particular, it can significantly improve the model’s ability to generate interrogative sentences that frequently appear in human-written review comments. This motivates future studies to consider building a large and more suitable pre-training dataset based on Stack Overflow and perform further pre-training if they aim to improve the review comment generation task.

We advocate future work to adopt Edit Progress to measure the partial progress achieved by ACR tools. The performances of code revision approaches have mostly been measured with a specific metric named Exact Match (EM). EM is very strict as it completely ignores progressive contributions (i.e., non-perfect results) of generated code by an approach. To measure the partial progressive aspect, we introduced Edit Progress (EP) and evaluated the approaches with such a metric. The overall results demonstrated that the rankings of models for each task could be changed according to which metric is being utilized. We study and verify the research value of such a metric here.

Figure 5 shows the EP scores of each test instance of CodeT5 and AutoTransform in the CRB task using Trans-Review_{data}. We choose these two approaches as they are the best-performing approaches for each metric on average in the CRB task. The vertical lines show the average EP scores of each approach. Note that we omit the instances whose EP scores are less than -50% for better visualization. As shown in Figure 5, CodeT5 has a better EM score because it has more generated perfect code predictions with 100% Edit Progress scores (identical to ground truths). From these vertical lines (i.e., the average EP scores), we confirm that AutoTransform, which is generally inferior to CodeT5 in EM, retrieves better EP scores. On the one hand, AutoTransform’s positive effects on the improvement of partial progress are demonstrated by

¹⁰Because both CodeT5 and T5-Review are pre-trained on the CodeSearchNet, to study the differences between them, we focus on analyzing Stack Overflow data for T5-Review.




Fig. 5. Edit Progress Scores of CodeT5 and AutoTransform in the Code Revision Before Review Task. The orange bars are the parts that AutoTransform surpasses CodeT5 in counts, and the blue bars are the parts that CodeT5 surpasses AutoTransform. The grey bars are the common parts in counts.

the more orange-colored bars in the high EP range (e.g., 80–99%). On the other hand, it generates much less noise than CodeT5, as we can verify with the figure (i.e., there are much fewer orange-colored bars under 0% EP). This phenomenon indicates that EP can evaluate the model performances in a more comprehensive way: not only focusing on those correct predictions but also taking the failure cases into account. EP can also evaluate the practical usability of ACR tools to some extent as users do not know which prediction is correct in advance and may expect all predictions have reasonably good qualities (e.g., at least making positive contributions).

B. Threats to Validity

General purpose pre-trained models diverge depending on different aspects, such as the characteristics of pre-training tasks, characteristics of downstream tasks, and the size of datasets. Our study may have a selection bias by considering several pre-trained models. We believe that employing three of the most popular pre-trained code models (i.e., CodeBERT, GraphCodeBERT, and CodeT5) can mitigate the threat. Another threat to validity regards the dataset selection, as it may deliver bias in the experimental results. In our work, we have minimized this threat by using multiple datasets, i.e., `Trans-Reviewdata` [20], `AutoTransformdata` [21], and `T5-Reviewdata` [22]. Furthermore, we publicly share the replication package¹¹ including datasets for future comparisons by the research community. The evaluation metrics used may also bring a selection bias. To reduce it, we reuse the same evaluation metric (i.e., Exact Match) with existing ACR tools [20], [21], [22] and the widely used metric for texts (i.e., BLEU [51]). Besides, to study the partial progress achieved

¹¹Replication Package. https://github.com/soarsmu/Generated_Review_How_Far

by ACR tools, we follow prior work in text editing [41] to adopt the Edit Progress metric.

VI. RELATED WORK

In this section, we review the research work that relates to our work: approaches to automating code review activities and pre-trained models for SE.

Automating Code Review Activities. We focus on three generation-based code review automation tasks: Code Revision Before Review [20], [21], [22], Code Revision After Review [20], [22], and Review Comment Generation [22]. Other studies focus on similar code review activities in different manners. Hellendoorn et al. [54] aimed to predict the locations of code change hunks that possibly need review and revision. They showed that locating a suitable hunk is a challenging task. Siow et al. [55] proposed CORE, a multi-level embedding approach to represent the semantics of code changes and reviews, which recommends code reviews in a retrieval-based way. Hong et al. [56] proposed a retrieval-based approach to recommend code review comments. Li et al. [57] developed CodeReviewer, a pre-trained model for code review. They evaluated it on three tasks using new datasets, aiming to improve code review. In contrast, we aimed to compare existing ACR tools and identify future directions.

In addition, there are other tasks to automate other code review activities to support software developers, such as reviewer recommendation [58], [59], [60], [61], review prioritization [62], [63], [64], [65] and defect-proneness prediction in submitted code [66], [67], [68], [69], [70], [71]. For instance, Thongtanunam et al. [58] found 4% to 30% of reviews in open-source projects could not find suitable reviewers (i.e., reviewer recommendation). Rahman et al. [59] proposed CORRECT which utilizes external library similarity and technology expertise similarity of reviewers to recommend reviewers. Al-Zubaidi et al. [61] did not only focus on reviewer experience but also take into account the review workload when recommending reviewers.

VII. CONCLUSION AND FUTURE WORK

In this paper, we empirically study the existing generation-based automatic code review (ACR) techniques and the general-purpose pre-trained code models. We evaluate the effectiveness of these tools on three representative generation-based code review tasks. Our experimental results show that CodeT5 is the best performer in most cases, outperforming the best existing ACR tool T5-Review by 13.4%–38.9% in Code Revision Before/After Review tasks. However, in Review Comment Generation, T5-Review outperforms all other methods by a large margin. Our work also has revealed several interesting findings that can motivate future work.

In the future, we plan to design a new ACR tool that can perform well both in Exact Match and Edit Progress metrics.

Acknowledgement. This research / project is supported by the National Research Foundation, Singapore, under its Industry Alignment Fund – Pre-positioning (IAF-PP) Funding

Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

REFERENCES

- [1] S. Wang, T. Liu, and L. Tan, "Automatically learning semantic features for defect prediction," in *Proceedings of the 38th International Conference on Software Engineering*, 2016, pp. 297–308.
- [2] R. K. Saha, M. Lease, S. Khurshid, and D. E. Perry, "Improving bug localization using structured information retrieval," in *2013 28th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 2013, pp. 345–355.
- [3] S. Wang and D. Lo, "Version history, similar report, and structure: Putting them together for improved bug localization," in *Proceedings of the 22nd International Conference on Program Comprehension*, 2014, pp. 53–63.
- [4] G. J. Myers, C. Sandler, and T. Badgett, *The art of software testing*. John Wiley & Sons, 2011.
- [5] V. Garousi and M. V. Mäntylä, "When and what to automate in software testing? a multi-vocal literature review," *Information and Software Technology*, vol. 76, pp. 92–117, 2016.
- [6] A. Bosu, J. C. Carver, C. Bird, J. Orbeck, and C. Chockley, "Process aspects and social dynamics of contemporary code review: Insights from open source development and industrial practice at microsoft," *IEEE Transactions on Software Engineering*, vol. 43, no. 1, pp. 56–75, 2016.
- [7] A. Bacchelli and C. Bird, "Expectations, outcomes, and challenges of modern code review," in *2013 35th International Conference on Software Engineering (ICSE)*. IEEE, 2013, pp. 712–721.
- [8] P. C. Rigby and C. Bird, "Convergent contemporary software peer review practices," in *Proceedings of the 2013 9th joint meeting on foundations of software engineering*, 2013, pp. 202–212.
- [9] S. McIntosh, Y. Kamei, B. Adams, and A. E. Hassan, "An empirical study of the impact of modern code review practices on software quality," *Empirical Software Engineering*, vol. 21, no. 5, pp. 2146–2189, 2016.
- [10] G. Bavota and B. Russo, "Four eyes are better than two: On the impact of code reviews on software quality," in *2015 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 2015, pp. 81–90.
- [11] R. Morales, S. McIntosh, and F. Khomh, "Do code review practices impact design quality? a case study of the qt, vtk, and itk projects," in *2015 IEEE 22nd international conference on software analysis, evolution, and reengineering (SANER)*. IEEE, 2015, pp. 171–180.
- [12] S. McIntosh, Y. Kamei, B. Adams, and A. E. Hassan, "The impact of code review coverage and code review participation on software quality: A case study of the qt, vtk, and itk projects," in *Proceedings of the 11th working conference on mining software repositories*, 2014, pp. 192–201.
- [13] C. Sadowski, E. Söderberg, L. Church, M. Sipko, and A. Bacchelli, "Modern code review: a case study at google," in *Proceedings of the 40th International Conference on Software Engineering: Software Engineering in Practice*, 2018, pp. 181–190.
- [14] M. Beller, A. Bacchelli, A. Zaidman, and E. Juergens, "Modern code reviews in open-source projects: Which problems do they fix?" in *Proceedings of the 11th working conference on mining software repositories*, 2014, pp. 202–211.
- [15] P. C. Rigby, D. M. German, L. Cowen, and M.-A. Storey, "Peer review on open-source software projects: Parameters, statistical models, and theory," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 23, no. 4, pp. 1–33, 2014.
- [16] A. J. Ko, B. A. Myers, M. J. Coblenz, and H. H. Aung, "An exploratory study of how developers seek, relate, and collect relevant information during software maintenance tasks," *IEEE Transactions on software engineering*, vol. 32, no. 12, pp. 971–987, 2006.
- [17] R. Minelli, A. Mocci, and M. Lanza, "I know what you did last summer—an investigation of how developers spend their time," in *2015 IEEE 23rd International Conference on Program Comprehension*. IEEE, 2015, pp. 25–35.
- [18] M. V. Zelkowitz, A. C. Shaw, and J. D. Gannon, *Principles of software engineering and design*. Prentice Hall Professional Technical Reference, 1979.
- [19] X. Xia, L. Bao, D. Lo, Z. Xing, A. E. Hassan, and S. Li, "Measuring program comprehension: A large-scale field study with professionals," *IEEE Transactions on Software Engineering*, vol. 44, no. 10, pp. 951–976, 2017.
- [20] R. Tufano, L. Pascarella, M. Tufano, D. Poshyvanyk, and G. Bavota, "Towards automating code review activities," in *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. IEEE, 2021, pp. 163–174.
- [21] P. Thongtanunam, C. Pornprasit, and C. Tantithamthavorn, "Autotransform: Automated code transformation to support modern code review process," in *2022 IEEE/ACM International Conference on Software Engineering (ICSE)*, 2022.
- [22] R. Tufano, S. Masiero, A. Mastropaolo, L. Pascarella, D. Poshyvanyk, and G. Bavota, "Using pre-trained models to boost code review automation," in *2022 IEEE/ACM International Conference on Software Engineering (ICSE)*, 2022.
- [23] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, Eds., 2017, pp. 5998–6008. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>
- [24] M. Tufano, J. Pantiuchina, C. Watson, G. Bavota, and D. Poshyvanyk, "On learning meaningful code changes via neural machine translation," in *Proceedings of the 41st International Conference on Software Engineering*, ser. ICSE '19, 2019.
- [25] R. Sennrich, B. Haddow, and A. Birch, "Neural machine translation of rare words with subword units," *arXiv preprint arXiv:1508.07909*, 2015.
- [26] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu *et al.*, "Exploring the limits of transfer learning with a unified text-to-text transformer," *J. Mach. Learn. Res.*, vol. 21, no. 140, pp. 1–67, 2020.
- [27] H. Husain, H. Wu, T. Gazit, M. Allamanis, and M. Brockschmidt, "Codesearchnet challenge: Evaluating the state of semantic code search," *CoRR*, vol. abs/1909.09436, 2019. [Online]. Available: <http://arxiv.org/abs/1909.09436>
- [28] Z. Feng, D. Guo, D. Tang, N. Duan, X. Feng, M. Gong, L. Shou, B. Qin, T. Liu, D. Jiang, and M. Zhou, "Codebert: A pre-trained model for programming and natural languages," in *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, ser. Findings of ACL, T. Cohn, Y. He, and Y. Liu, Eds., vol. EMNLP 2020. Association for Computational Linguistics, 2020, pp. 1536–1547. [Online]. Available: <https://doi.org/10.18653/v1/2020.findings-emnlp.139>
- [29] D. Guo, S. Ren, S. Lu, Z. Feng, D. Tang, S. Liu, L. Zhou, N. Duan, A. Svyatkovskiy, S. Fu, M. Tufano, S. K. Deng, C. B. Clement, D. Drain, N. Sundaresan, J. Yin, D. Jiang, and M. Zhou, "Graphcodebert: Pre-training code representations with data flow," in *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [Online]. Available: <https://openreview.net/forum?id=jLoC4ez43PZ>
- [30] Y. Wang, W. Wang, S. R. Joty, and S. C. H. Hoi, "Codet5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, M. Moens, X. Huang, L. Specia, and S. W. Yih, Eds. Association for Computational Linguistics, 2021, pp. 8696–8708. [Online]. Available: <https://doi.org/10.18653/v1/2021.emnlp-main.685>
- [31] X. Zhou, D. Han, and D. Lo, "Assessing generalizability of codebert," in *2021 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 2021, pp. 425–436.
- [32] S. Lu, D. Guo, S. Ren, J. Huang, A. Svyatkovskiy, A. Blanco, C. Clement, D. Drain, D. Jiang, D. Tang *et al.*, "Codexglue: A machine learning benchmark dataset for code understanding and generation," *arXiv preprint arXiv:2102.04664*, 2021.
- [33] E. Winter, D. Bowes, S. Counsell, T. Hall, S. Haraldsson, V. Nowack, and J. Woodward, "How do developers really feel about bug fixing? directions for automatic program repair," *IEEE Transactions on Software Engineering*, 2022.
- [34] P. S. Kochhar, F. Thung, N. Nagappan, T. Zimmermann, and D. Lo, "Understanding the test automation culture of app developers," in *2015 IEEE 8th International Conference on Software Testing, Verification and Validation (ICST)*. IEEE, 2015, pp. 1–10.

- [35] P. S. Kochhar, X. Xia, D. Lo, and S. Li, "Practitioners' expectations on automated fault localization," in *Proceedings of the 25th International Symposium on Software Testing and Analysis*, 2016, pp. 165–176.
- [36] D. Linsley, J. Kim, V. Veerabadran, C. Windolf, and T. Serre, "Learning long-range spatial dependencies with horizontal gated recurrent units," *Advances in neural information processing systems*, vol. 31, 2018.
- [37] V. J. Hellendoorn and P. Devanbu, "Are deep neural networks the best choice for modeling source code?" in *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*, 2017, pp. 763–773.
- [38] R.-M. Karampatsis, H. Babii, R. Robbes, C. Sutton, and A. Janes, "Big code!= big vocabulary: Open-vocabulary models for source code," in *2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE)*. IEEE, 2020, pp. 1073–1085.
- [39] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, J. Burstein, C. Doran, and T. Solorio, Eds. Association for Computational Linguistics, 2019, pp. 4171–4186. [Online]. Available: <https://doi.org/10.18653/v1/n19-1423>
- [40] K. Clark, M. Luong, Q. V. Le, and C. D. Manning, "ELECTRA: pre-training text encoders as discriminators rather than generators," in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. [Online]. Available: <https://openreview.net/forum?id=r1xMH1BtvB>
- [41] A. Elgohary, C. Meek, M. Richardson, A. Fournay, G. Ramos, and A. H. Awadallah, "NL-EDIT: Correcting semantic parse errors through natural language interaction," in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Online: Association for Computational Linguistics, Jun. 2021, pp. 5599–5610. [Online]. Available: <https://aclanthology.org/2021.naacl-main.444>
- [42] V. Dibia, A. Fournay, G. Bansal, F. Poursabzi-Sangdeh, H. Liu, and S. Amershi, "Aligning offline metrics and human judgments of value of ai-pair programmers," *arXiv preprint arXiv:2210.16494*, 2022.
- [43] 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.
- [44] V. I. Levenshtein *et al.*, "Binary codes capable of correcting deletions, insertions, and reversals," in *Soviet physics doklady*, vol. 10, no. 8. Soviet Union, 1966, pp. 707–710.
- [45] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann *et al.*, "Palm: Scaling language modeling with pathways," *arXiv preprint arXiv:2204.02311*, 2022.
- [46] W. Ma, Y. Cui, C. Si, T. Liu, S. Wang, and G. Hu, "Charbert: character-aware pre-trained language model," *arXiv preprint arXiv:2011.01513*, 2020.
- [47] L. Xue, A. Barua, N. Constant, R. Al-Rfou, S. Narang, M. Kale, A. Roberts, and C. Raffel, "Byt5: Towards a token-free future with pre-trained byte-to-byte models," *Transactions of the Association for Computational Linguistics*, vol. 10, pp. 291–306, 2022.
- [48] S. J. Mielke, Z. Alyafeai, E. Salesky, C. Raffel, M. Dey, M. Gallé, A. Raja, C. Si, W. Y. Lee, B. Sagot *et al.*, "Between words and characters: A brief history of open-vocabulary modeling and tokenization in nlp," *arXiv preprint arXiv:2112.10508*, 2021.
- [49] J. P. Turian, L. Shea, and I. D. Melamed, "Evaluation of machine translation and its evaluation," in *Machine Translation Summit*, 2003.
- [50] D. A. Coughlin, "Correlating automated and human assessments of machine translation quality," in *Machine Translation Summit*, 2003.
- [51] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "Bleu: a method for automatic evaluation of machine translation," in *Annual Meeting of the Association for Computational Linguistics*, 2002.
- [52] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [53] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, "Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension," *arXiv preprint arXiv:1910.13461*, 2019.
- [54] V. J. Hellendoorn, J. Tsay, M. Mukherjee, and M. Hirzel, "Towards automating code review at scale," in *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2021, pp. 1479–1482.
- [55] J. K. Siow, C. Gao, L. Fan, S. Chen, and Y. Liu, "Core: Automating review recommendation for code changes," in *2020 IEEE 27th International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 2020, pp. 284–295.
- [56] Y. Hong, C. Tantithamthavorn, P. Thongtanunam, and A. Aleti, "Commentfinder: a simpler, faster, more accurate code review comments recommendation," in *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2022, pp. 507–519.
- [57] Z. Li, S. Lu, D. Guo, N. Duan, S. Jannu, G. Jenks, D. Majumder, J. Green, A. Svyatkovskiy, S. Fu *et al.*, "Automating code review activities by large-scale pre-training," in *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2022, pp. 1035–1047.
- [58] P. Thongtanunam, C. Tantithamthavorn, R. G. Kula, N. Yoshida, H. Iida, and K.-i. Matsumoto, "Who should review my code? a file location-based code-reviewer recommendation approach for modern code review," in *2015 IEEE 22nd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*. IEEE, 2015, pp. 141–150.
- [59] M. M. Rahman, C. K. Roy, J. Redl, and J. A. Collins, "Correct: Code reviewer recommendation at github for vendasta technologies," in *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering*, 2016, pp. 792–797.
- [60] E. Sülün, E. Tüzün, and U. Doğrusöz, "Reviewer recommendation using software artifact traceability graphs," in *Proceedings of the fifteenth international conference on predictive models and data analytics in software engineering*, 2019, pp. 66–75.
- [61] W. H. A. Al-Zubaidi, P. Thongtanunam, H. K. Dam, C. Tantithamthavorn, and A. Ghose, "Workload-aware reviewer recommendation using a multi-objective search-based approach," in *Proceedings of the 16th ACM International Conference on Predictive Models and Data Analytics in Software Engineering*, 2020, pp. 21–30.
- [62] Y. Fan, X. Xia, D. Lo, and S. Li, "Early prediction of merged code changes to prioritize reviewing tasks," *Empirical Software Engineering*, vol. 23, no. 6, pp. 3346–3393, 2018.
- [63] S.-T. Shi, M. Li, D. Lo, F. Thung, and X. Huo, "Automatic code review by learning the revision of source code," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 4910–4917.
- [64] H.-Y. Li, S.-T. Shi, F. Thung, X. Huo, B. Xu, M. Li, and D. Lo, "Deep-review: automatic code review using deep multi-instance learning," in *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 2019, pp. 318–330.
- [65] B. Wu, B. Liang, and X. Zhang, "Turn tree into graph: Automatic code review via simplified ast driven graph convolutional network," *Knowledge-Based Systems*, vol. 252, p. 109450, 2022.
- [66] Y. Kamei, E. Shihab, B. Adams, A. E. Hassan, A. Mockus, A. Sinha, and N. Ubayashi, "A large-scale empirical study of just-in-time quality assurance," *IEEE Transactions on Software Engineering*, vol. 39, no. 6, pp. 757–773, 2012.
- [67] S. McIntosh and Y. Kamei, "Are fix-inducing changes a moving target? a longitudinal case study of just-in-time defect prediction," in *Proceedings of the 40th International Conference on Software Engineering*, 2018, pp. 560–560.
- [68] T. Hoang, H. K. Dam, Y. Kamei, D. Lo, and N. Ubayashi, "Deepjit: an end-to-end deep learning framework for just-in-time defect prediction," in *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*. IEEE, 2019, pp. 34–45.
- [69] T. Hoang, H. J. Kang, D. Lo, and J. Lawall, "Cc2vec: distributed representations of code changes," in *ICSE '20: 42nd International Conference on Software Engineering, Seoul, South Korea, 27 June - 19 July, 2020*, G. Rothermel and D. Bae, Eds. ACM, 2020, pp. 518–529. [Online]. Available: <https://doi.org/10.1145/3377811.3380361>
- [70] C. Pornprasit and C. K. Tantithamthavorn, "Jitline: A simpler, better, faster, finer-grained just-in-time defect prediction," in *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. IEEE, 2021, pp. 369–379.
- [71] X. Zhou, D. Han, and D. Lo, "Simple or complex? together for a more accurate just-in-time defect predictor," in *2022 IEEE/ACM 30th International Conference on Program Comprehension (ICPC)*. IEEE, 2022, pp. 229–240.