

# Enhancing Code Annotation Reliability: Generative AI's Role in Comment Quality Assessment Models

Seetharam Killivalavan<sup>1,\*</sup>, Durairaj Thenmozhi<sup>2,\*</sup>

<sup>1</sup>Sri Sivasubramaniya Nadar College of Engineering, Chennai, Tamil Nadu- 603110

## Abstract

This paper explores a novel method for enhancing binary classification models that assess code comment quality, leveraging Generative Artificial Intelligence to elevate model performance. By integrating 1,437 newly generated code-comment pairs, labeled as "Useful" or "Not Useful" and sourced from various GitHub repositories, into an existing C-language dataset of 9,048 pairs, we demonstrate substantial model improvements. Using an advanced Large Language Model, our approach yields a 5.78% precision increase in the Support Vector Machine (SVM) model, improving from 0.79 to 0.8478, and a 2.17% recall boost in the Artificial Neural Network (ANN) model, rising from 0.731 to 0.7527. These results underscore Generative AI's value in advancing code comment classification models, offering significant potential for enhanced accuracy in software development and quality control. This study provides a promising outlook on the integration of generative techniques for refining machine learning models in practical software engineering settings.

## Keywords

Code Comment Quality Classification, Generative Artificial Intelligence, Support Vector Machines, Artificial Neural Networks, Natural Language Processing

## 1. Introduction

Code comments are essential in software development, enhancing understanding, supporting team collaboration, and facilitating long-term code maintenance, as discussed by De et al. (2005) [1]. However, manually evaluating comments poses challenges due to its time-intensive and subjective nature, as noted by Haouari et al. (2011) [2]. To address these limitations, this study explores the use of Generative AI to automate comment quality assessment, as proposed by Ebert et al. (2023) [3], presenting a significant advancement for optimizing code review processes and expediting the Software Development Life Cycle (SDLC).

Incorporating comments effectively within the SDLC can benefit developers by accelerating troubleshooting, providing essential documentation, and establishing a robust groundwork for future development phases, as suggested by Majumdar (2020) [4]. This paper details our methods, experimental design, and the transformative potential of this AI-based approach for the software engineering field, as previously highlighted by Roehm et al. (2012) [5]. Following this introduction, we review existing studies on comment classification and explain our process for generating a new dataset using Large Language Models (LLMs).

### 1.1. CODE COMMENT CLASSIFICATION: CURRENT LANDSCAPE AND CHALLENGES

Code comments are used to clarify logic, design decisions, and develop challenges [6]. However, manual evaluation remains inconsistent, time-consuming, and subjective [4]. Automated classification, labeling comments as "Useful" or "Not Useful," offers a more efficient approach to streamline code review [7]. This study examines how Generative AI can enhance these classification models [3], potentially transforming comment quality assessment. By prioritizing essential comments, resource management can improve.

Forum for Information Retrieval Evaluation, December 12-15, 2024, India

\*Corresponding author.

✉ seetharam2210463@ssn.edu.in (S. Killivalavan); theni\$.d\$@ssn.edu.in (D. Thenmozhi)

🆔 0009-0002-2655-117X (S. Killivalavan); 0000-0003-0681-6628 (D. Thenmozhi)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0)

This introduction sets up a discussion on how Large Language Models (LLMs) are advancing code comment classification and software development practices [1].

## 1.2. IMPACT OF LLM ON THE QUALITY OF COMMENTS

Leveraging Large Language Models (LLMs) represents a major advancement in evaluating the quality of code comments [3]. These models move beyond syntactic comprehension, capturing the deeper semantics of the code and generating insightful comments that streamline assessment processes. By doing so, they significantly enhance the relevance and clarity of comments across the Software Development Life Cycle (SDLC). Beyond mere classification, LLMs redefine developer interaction with code, fostering clearer communication and strengthening collaboration. This transformative impact underscores the essential role LLMs are set to play in the future of code comment quality evaluation.

The application of Generative AI within the IRSE@FIRE-2024 task [8] is set to transform code quality evaluation, streamlining the Software Development Life Cycle (SDLC) and promoting more effective resource distribution and collaborative development efforts among teams.

The subsequent sections are organized as follows: Section 2 provides an overview of comment classification and the foundations of Generative AI. Section 3 describes the task setup and dataset used. Our methodology is detailed in Section 4. In Section 5, we present the results, while Section 6 offers a comparative analysis of our models and embeddings against established approaches in code comment quality assessment, underscoring their unique contributions. Lastly, Section 7 concludes with a summary of our findings and discusses possible avenues for future research.

## 2. Related Work

Automated program understanding is a recognized research area among professionals in the software domain. Various tools have been developed to facilitate the extraction of knowledge from software metadata, encompassing components such as runtime traces and structural attributes of code [1, 9, 10, 11, 12, 13, 14, 15]. Researchers have developed various methods to mine and evaluate code comments, focusing on analyzing comment quality through code-comment pair comparisons. In assessing code comment quality, authors [16, 17, 18, 19, 20, 21, 22, 23] employ techniques such as word similarity measures (e.g., Levenshtein distance) and comment length analysis to filter out trivial and non-informative comments. Rahman et al. [24] detect useful and non-useful code review comments (logged-in review portals) based on attributes identified from a survey conducted with developers of Microsoft [25].

New programmers often rely on existing comments to comprehend code flow. However, not all comments contribute effectively to program comprehension, necessitating a relevancy assessment of source code comments prior to their use. Numerous researchers have focused on the automatic classification of source code comments in terms of quality evaluation. For instance, Omal et al. [26] noted that factors influencing software maintainability can be organized into hierarchical structures. The authors defined measurable attributes in the form of metrics for each factor, enabling the assessment of software characteristics, which can then be consolidated into a single index of software maintainability. Fluri et al. [27] examined whether the source code and associated comments are changed together along the multiple versions. They investigated three open source systems, such as *ArgoUML*, *Azureus*, and *JDT Core*, and found that 97% of the comment changes are done in the same revision as the associated source code changes. Yu Hai et al. [28] classified source code comments into four classes - unqualified, qualified, good, and excellent. The aggregation of basic classification algorithms further improved the classification result. Another work published in [7] in which author proposed an automatic classification mechanism "CommentProbe" for quality evaluation of code comments of C codebases. We see that people worked on source code comments with different aspects [7, 4, 20, 19, 22, 23], but still, automatic quality evaluation of source code comments is an important area and demands more research.

With the advent of large language models [29], it is important to compare the quality assessment of code comments by the standard models like GPT 3.5 or llama with the human interpretation. The IRSE track at FIRE 2024 [30, 31] builds upon the methodologies proposed in [7, 32, 8, 19] to investigate various vector space models [33] and features for binary classification and evaluation of comments in relation to code comprehension. This track also assesses the performance of the predictive model by incorporating GPT-generated labels for the quality of code and comment snippets extracted from open-source software.

### 3. Task and Dataset Description

This section outlines the IRSE@FIRE-2024 task [8], focused on improving a binary code comment quality classification model. The task involves integrating newly generated code-comment pairs for enhanced accuracy. It comprises an initial dataset of 9048 labeled code-comment pairs in C, out of which 5378 were classified as "Useful" and 3670 were classified as "Not Useful", along with additional pairs generated using a Large Language Model (LLM), each labeled.

The desired output includes two versions of the classification model: one with the added generated pairs and labels, and another without. The starting dataset encompasses 9048 comments from GitHub, each with the comment text, surrounding code, and a corresponding usefulness label (Table 1).

**Table 1**  
Sample Data Instance

Comment	Code	Label
/* Swap two values */	void swapValues(int *x, int *y) { int temp; temp = *x; *x = *y; *y = temp;}	Useful
/* Compute the sixth power */	int result = computeSixth-Power(value);	Not Useful
/*Simple variable declaration */	int x = 10;	Not Useful

To establish the ground truth, 14 annotators assessed each comment independently, resulting in substantial agreement (Cohen's kappa value of 0.734). The annotation process involved the assessment of a comprehensive set of 16,000 comments.

Participants are also tasked with generating an additional dataset of labeled code-comment pairs from GitHub using an LLM. This dataset is to be submitted alongside the task.

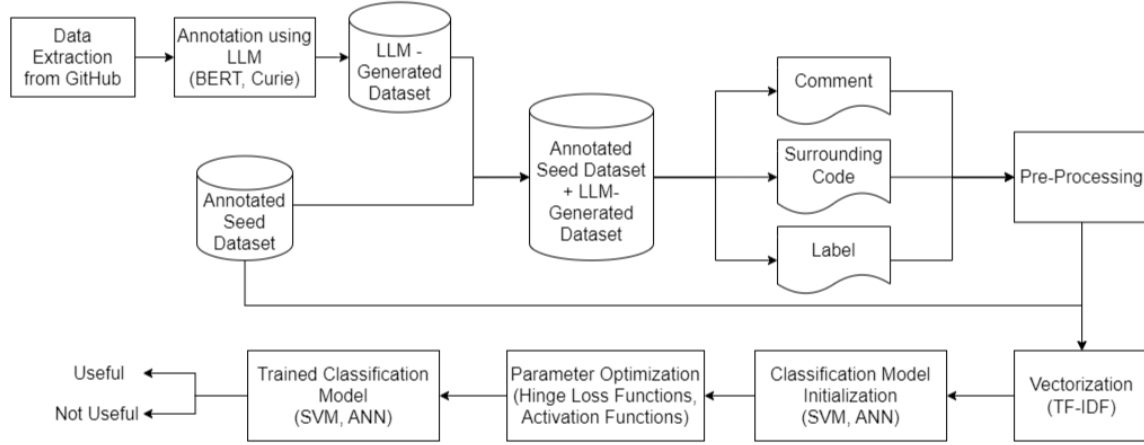
In summary, the objective is to refine the code comment quality classification model by integrating newly generated pairs, ultimately enhancing accuracy and effectiveness.

For further details, please refer to the task description provided at IRSE@FIRE-2024 <sup>1</sup>.

### 4. Methodology

Our approach encompasses the combination robust methodologies, including Support Vector Machine (SVM) models for classification and Artificial Neural Networks (ANN) with diverse activation functions for capturing complex data relationships [34]. Additionally, we leverage Large Language Models (LLMs) via the OpenAI API and utilize GitHub repositories to generate a diverse and substantial dataset of code-comment pairs. The following subtopics detail our specific methodologies: implementing SVM models, exploring ANN models, and generating datasets using the OpenAI API and GitHub repositories. These methodologies collectively form the foundation of our innovative approach to code comment quality assessment. Within the framework of our methodology, Figure 1 elegantly elucidates the architectural blueprint that underpins our approach.

<sup>1</sup><https://sites.google.com/view/irse2024/home>



**Figure 1:** Architecture Diagram

## 4.1. Support Vector Machines

A Linear Support Vector Machine (SVM) is a powerful classification technique that finds the optimal hyperplane for effective data separation, expressed as  $y = mx + b$ , where  $y$  is the predicted class label,  $x$  is the input data,  $m$  is the slope and  $b$  is the y-intercept. It maximizes the margin, which is the distance between the hyperplane and the nearest data points. This margin ( $M$ ) can be calculated as:

$$M = \frac{2}{\|m\|} \quad (1)$$

where  $\|m\|$  is the length of the weight vector  $m$ .

SVM aims to minimize the square of the length of the weight vector ( $\|m\|^2$ ) while ensuring that each data point  $x_i$  is correctly classified:

$$y_i(m \cdot x_i + b) \geq 1 \quad (2)$$

Equation 2 states that the product must be greater than or equal to 1 for all data points, emphasizing the importance of well-defined class separation in SVM classification. This condition is central to SVM's goal of locating an optimal hyperplane, maximizing the margin, and guaranteeing accurate data point classification. Support vectors, those closest to the hyperplane, are pivotal in margin definition, thereby influencing SVM's overall performance.

## 4.2. Artificial Neural Networks

Artificial Neural Networks (ANNs) are adaptable machine learning models that draw inspiration from the architecture and operation of the human brain. They excel at discerning complex data relationships, making them highly effective for tasks like code comment quality classification. The mathematical representation of a single neuron in an ANN is given by:

$$Z = w_1x_1 + w_2x_2 + \dots + w_nx_n + b \quad (3)$$

where  $x_n$  are input features,  $w_n$  are corresponding weights and  $b$  is the bias term.

The weighted sum ( $Z$ ) is then passed through an activation function, which introduces non-linearity into the model. Different activation functions yield different learning behaviours.

Here are a few common activation functions and their formulas:

i) Logistic Function:

$$f(Z) = \frac{1}{1 + e^{-Z}} \quad (4)$$

ii) Rectified Linear Unit (ReLU):

$$f(Z) = \max(0, Z) \quad (5)$$

iii) Hyperbolic Tangent (tanh):

$$f(Z) = \frac{e^Z - e^{-Z}}{e^Z + e^{-Z}} \quad (6)$$

### 4.3. Leveraging LLM for Generation of Dataset

Our methodology encompasses a multi-step approach to dataset generation. Initially, we leveraged both the OpenAI API, powered by the Curie Model, and GitHub repositories to diversify our dataset. The API simulated real-world coding scenarios, producing authentic code-comment pairs and substantially augmenting our dataset. Complementing this, we extracted additional pairs from various open-source projects on GitHub, ensuring relevance and utility. This combined strategy significantly broadened the dataset's coverage while upholding high quality standards. Subsequently, the code-comment pairs underwent processing using OpenAI's Curie Model in conjunction with BERT for label generation, signifying comment usefulness. This involved presenting prompts with both code and comment, and employing the LLM to generate a label. Finally, the dataset was meticulously assembled, each entry comprising code, comment, and the corresponding generated label. This rigorous methodology serves as a robust foundation for our code comment quality classification model.

## 5. Analysis of Results

Evaluating our code comment quality classification model is a crucial step in validating its effectiveness. We utilized a combination of Support Vector Machines (SVM) and Artificial Neural Networks (ANN) with various activation functions, including ReLU, identity, logistic, and tanh, to conduct a comprehensive analysis of the model's performance. This multidimensional approach offered valuable insights into the model's adaptability, revealing its robustness across diverse scenarios. Additionally, integrating these methodologies resulted in a significant improvement in precision, underscoring the model's ability to categorize code comments accurately based on practical value. These findings align with previous research that demonstrates the reliability of SVM and ANN models for comment quality assessment. The use of diverse activation functions further highlights the flexibility of our approach, reinforcing the model's potential applicability in real-world software development.

### 5.1. Classification Models

The evaluation of our code comment quality classification models yielded insightful findings, showcasing the impact of integrating LLM-generated data into our seed dataset of 9048 entries. This initial dataset was thoughtfully partitioned into training, testing and validation sets, with the testing set comprising 1718 entries. With the Seed Data, SVM exhibited commendable precision (0.79), while ANN with ReLU activation demonstrated remarkable effectiveness, resulting in a notable recall score (0.731). Models with tanh and logistic activation functions showed similar precision scores of 0.726 and 0.73.

Post integration of 1437 LLM-generated entries, which seamlessly enriched the Seed Data, SVM's precision notably increased by 5.78%, elevating the preceding value to 0.8478, highlighting the value of incorporating generative AI. Using ReLU, ANN achieved a noteworthy 2.17% rise in its recall, giving it a final recall of 0.7527, while tanh and logistic functions yielded marginal changes. Extensive experimentation with varied SVM models and ANN activation functions was performed, and the results depicts the effectiveness of our approach, emphasizing the importance of meticulous experimentation in fine-tuning models for code comment quality analysis.

Furthermore, for detailed numerical insights, please refer to Table 2, which provides a comparison of the model performance, offering the classification report of our top-performing models. It serves as a comprehensive reference for our findings and allows the comparison of test accuracies and F1 scores before and after integration.

**Table 2**  
Model Performance Comparison

Model	Performance with Seed Data		Performance with Integrated Data	
	Test Accuracy	F1-Score	Test Accuracy	F1-Score
Linear SVM	0.811	0.792	0.813	0.79
SVM (poly. kernel)	0.798	0.781	0.814	0.796
ANN (ReLU)	0.75	0.741	0.757	0.744
ANN (tanh)	0.753	0.729	0.749	0.743
ANN (logistic)	0.749	0.73	0.745	0.736
ANN (identity)	0.748	0.728	0.743	0.729

## 5.2. Analysis of Dataset Generated using LLM

The integration of data generated by OpenAI's Large Language Model (LLM), in conjunction with the utilization of the Curie model, and the inclusion of diverse datasets from various GitHub repositories and open-source projects represents a significant stride in elevating our code comment quality classification model. By meticulously adding 1437 new entries to our original dataset, we substantially enriched the diversity of our training corpus. This augmentation in data diversity led to a marked improvement in the accuracy of our classification model, benefiting both Support Vector Machine (SVM) and Artificial Neural Network (ANN) models. The heightened sensitivity achieved through this amalgamation enhances the model's generalization and prediction capabilities, underscoring the value of incorporating external data sources. Furthermore, the integration of BERT embeddings and the Curie model empowered our model to adeptly capture the intricacies of code commentary, notably enhancing its ability to distinguish between "Useful" and "Not Useful" comments. This capability proves crucial in real-world scenarios, where precise comment assessment plays a pivotal role in influencing the effectiveness of software development and maintenance processes.

## 6. Discussion

In this section, we conduct a thorough comparative analysis of our models and embeddings in relation to previous studies on code comment classification. Our deliberate emphasis on Support Vector Machine (SVM) and Artificial Neural Network (ANN) models, each with specific activation functions, allows for an in-depth exploration of their efficacy. This focused investigation provides nuanced insights into their performance in code comment quality assessment, contrasting with the broader set of classifiers utilized by Majumdar et al. (2022a) [7].

Additionally, our research methodology diverges from the work of Majumdar et al. (2020) [4], which primarily centers on the extraction of knowledge domains from code comments for addressing developer queries during maintenance. In contrast, our focus centers on the development and evaluation of code comment quality classification models. This includes the integration of LLM-generated data, resulting in significant enhancements in classification precision.

Concerning embeddings, Majumdar et al. (2022b) [33] emphasize contextualized word representations fine-tuned on software development texts. In our case, we utilized both BERT and custom embeddings specifically tailored for software development concepts. This approach provided high-dimensional semantic representations, catering to a wide array of natural language processing tasks. It's worth noting that for labeling, we harnessed the Curie model. This distinction underscores the versatility



and broader applicability of our embeddings compared to the contextualized embeddings discussed by Majumdar et al (2022b)[33].

Fundamentally, our proposition emphatically focuses on specific models and embeddings, providing unique insights into their effectiveness for assessing code comment quality. The emphasis on specific models and customized embeddings offers detailed insights into evaluating code comment quality, distinguishing it from the broader, contextually-focused techniques utilized in prior research [33].

## 7. Conclusion

Building on these foundational advancements, our study highlights the practicality and scalability of Generative AI for real-world applications. By generating and integrating new data into existing datasets, we demonstrated that Generative AI could enhance the performance of traditional models in code comment quality assessment. This approach not only elevated our models' precision and recall but also underscored the potential of Generative AI to provide robust solutions for improving software documentation practices, making it an impactful tool for future development cycles.

The integration of LLM-generated data notably amplified model performance, with precision for the SVM model increasing by 5.78% and recall for the ANN model improving by 2.17%. These enhancements raised the test accuracies to 81.1% for SVM and 75% for ANN, marking a clear advancement from their pre-augmentation baselines. These quantifiable gains underscore the effectiveness of data augmentation via Generative AI, illustrating how even modest dataset expansions can yield substantial improvements in model accuracy and reliability, particularly for complex classification tasks in software development.

Looking ahead, the impact of this work can extend well beyond code comment classification. The methodologies introduced here establish a versatile framework that can be adapted for a wide range of tasks in software development and quality assurance. By leveraging generative AI, specifically through Large Language Models (LLMs), we highlight a powerful approach that could redefine code analysis and documentation practices. As the software industry evolves, this research stands as evidence of the substantial value in adopting advanced technologies, reinforcing the importance of innovative solutions in enhancing efficiency and precision in practical engineering applications.

## Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT in order to: Grammar and spelling check. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

## References

- [1] S. C. B. de Souza, N. Anquetil, K. M. de Oliveira, A study of the documentation essential to software maintenance, Conference on Design of communication, ACM, 2005, pp. 68–75.
- [2] D. Haouari, H. Sahraoui, P. Langlais, How good is your comment? a study of comments in java programs, in: 2011 International symposium on empirical software engineering and measurement, IEEE, 2011, pp. 137–146.
- [3] C. Ebert, P. Louridas, Generative ai for software practitioners, IEEE Software 40 (2023) 30–38.
- [4] S. Majumdar, S. Papdeja, P. P. Das, S. K. Ghosh, Comment-mine—a semantic search approach to program comprehension from code comments, in: Advanced Computing and Systems for Security, Springer, 2020, pp. 29–42.
- [5] T. Roehm, R. Tiarks, R. Koschke, W. Maalej, How do professional developers comprehend software?, in: 2012 34th International Conference on Software Engineering (ICSE), IEEE, 2012, pp. 255–265.

- [6] P. Rani, S. Panichella, M. Leuenberger, A. Di Sorbo, O. Nierstrasz, How to identify class comment types? a multi-language approach for class comment classification, *Journal of systems and software* 181 (2021) 111047.
- [7] S. Majumdar, A. Bansal, P. P. Das, P. D. Clough, K. Datta, S. K. Ghosh, Automated evaluation of comments to aid software maintenance, *Journal of Software: Evolution and Process* 34 (2022) e2463.
- [8] S. Majumdar, S. Paul, D. Paul, A. Bandyopadhyay, S. Chattopadhyay, P. P. Das, P. D. Clough, P. Majumder, Generative ai for software metadata: Overview of the information retrieval in software engineering track at fire 2023, *arXiv preprint arXiv:2311.03374* (2023).
- [9] S. Majumdar, S. Papdeja, P. P. Das, S. K. Ghosh, Smartkt: a search framework to assist program comprehension using smart knowledge transfer, in: *2019 IEEE 19th International Conference on Software Quality, Reliability and Security (QRS)*, IEEE, 2019, pp. 97–108.
- [10] N. Chatterjee, S. Majumdar, S. R. Sahoo, P. P. Das, Debugging multi-threaded applications using pin-augmented gdb (pgdb), in: *International conference on software engineering research and practice (SERP)*. Springer, 2015, pp. 109–115.
- [11] S. Majumdar, N. Chatterjee, S. R. Sahoo, P. P. Das, D-cube: tool for dynamic design discovery from multi-threaded applications using pin, in: *2016 IEEE International Conference on Software Quality, Reliability and Security (QRS)*, IEEE, 2016, pp. 25–32.
- [12] S. Majumdar, N. Chatterjee, P. P. Das, A. Chakrabarti, A mathematical framework for design discovery from multi-threaded applications using neural sequence solvers, *Innovations in Systems and Software Engineering* 17 (2021) 289–307.
- [13] S. Majumdar, N. Chatterjee, P. Pratim Das, A. Chakrabarti, Dcube\_ nn d cube nn: Tool for dynamic design discovery from multi-threaded applications using neural sequence models, *Advanced Computing and Systems for Security: Volume 14* (2021) 75–92.
- [14] J. Siegmund, N. Peitek, C. Parnin, S. Apel, J. Hofmeister, C. Kästner, A. Begel, A. Bethmann, A. Brechmann, Measuring neural efficiency of program comprehension, in: *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*, 2017, pp. 140–150.
- [15] N. Chatterjee, S. Majumdar, P. P. Das, A. Chakrabarti, Parallelc-assist: Productivity accelerator suite based on dynamic instrumentation, *IEEE Access* (2023).
- [16] L. Tan, D. Yuan, Y. Zhou, Hotcomments: how to make program comments more useful?, in: *Conference on Programming language design and implementation (SIGPLAN)*, ACM, 2007, pp. 20–27.
- [17] Y. Wang, H. Le, A. D. Gotmare, N. D. Bui, J. Li, S. C. Hoi, Codet5+: Open code large language models for code understanding and generation, *arXiv preprint arXiv:2305.07922* (????).
- [18] D. Steidl, B. Hummel, E. Juergens, Quality analysis of source code comments, *International Conference on Program Comprehension (ICPC)*, IEEE, 2013, pp. 83–92.
- [19] S. Majumdar, A. Bandyopadhyay, P. P. Das, P. Clough, S. Chattopadhyay, P. Majumder, Can we predict useful comments in source codes?-analysis of findings from information retrieval in software engineering track@ fire 2022, in: *Proceedings of the 14th Annual Meeting of the Forum for Information Retrieval Evaluation*, 2022, pp. 15–17.
- [20] S. Majumdar, A. Bandyopadhyay, S. Chattopadhyay, P. P. Das, P. D. Clough, P. Majumder, Overview of the irse track at fire 2022: Information retrieval in software engineering., in: *FIRE (Working Notes)*, 2022, pp. 1–9.
- [21] J. L. Freitas, D. da Cruz, P. R. Henriques, A comment analysis approach for program comprehension, *Annual Software Engineering Workshop (SEW)*, IEEE, 2012, pp. 11–20.
- [22] S. Majumdar, P. P. Das, Smart knowledge transfer using google-like search, *arXiv preprint arXiv:2308.06653* (2023).
- [23] P. Chakraborty, S. Dutta, D. K. Sanyal, S. Majumdar, P. P. Das, Bringing order to chaos: Conceptualizing a personal research knowledge graph for scientists., *IEEE Data Eng. Bull.* 46 (2023) 43–56.
- [24] M. M. Rahman, C. K. Roy, R. G. Kula, Predicting usefulness of code review comments using textual features and developer experience, *International Conference on Mining Software Repositories*



- (MSR), IEEE, 2017, pp. 215–226.
- [25] A. Bosu, M. Greiler, C. Bird, Characteristics of useful code reviews: An empirical study at microsoft, Working Conference on Mining Software Repositories, IEEE, 2015, pp. 146–156.
  - [26] P. Oman, J. Hagemester, Metrics for assessing a software system’s maintainability, in: Proceedings Conference on Software Maintenance 1992, IEEE Computer Society, 1992, pp. 337–338.
  - [27] B. Fluri, M. Wursch, H. C. Gall, Do code and comments co-evolve? on the relation between source code and comment changes, in: 14th Working Conference on Reverse Engineering (WCRE 2007), IEEE, 2007, pp. 70–79.
  - [28] H. Yu, B. Li, P. Wang, D. Jia, Y. Wang, Source code comments quality assessment method based on aggregation of classification algorithms, Journal of Computer Applications 36 (2016) 3448.
  - [29] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, Advances in neural information processing systems 33 (2020) 1877–1901.
  - [30] S. Paul, S. Majumdar, R. Shah, S. Das, M. Ghosh, D. Ganguly, G. Calikli, D. Sanyal, P. P. Das, P. D Clough, A. Bandyopadhyay, S. Chattopadhyay, Generative ai for code metadata quality assessment, in: Proceedings of the 16th Annual Meeting of the Forum for Information Retrieval Evaluation, 2024.
  - [31] S. Paul, S. Majumdar, R. Shah, S. Das, M. Ghosh, D. Ganguly, G. Calikli, D. Sanyal, P. P. Das, P. D Clough, A. Bandyopadhyay, S. Chattopadhyay, Overview of the irse track at fire 2024: Information retrieval in software engineering, in: FIRE (Working Notes), 2024.
  - [32] S. Paul, S. Majumdar, A. Bandyopadhyay, B. Dave, S. Chattopadhyay, P. Das, P. D. Clough, P. Majumder, Efficiency of large language models to scale up ground truth: Overview of the irse track at forum for information retrieval 2023, in: Proceedings of the 15th Annual Meeting of the Forum for Information Retrieval Evaluation, 2023, pp. 16–18.
  - [33] S. Majumdar, A. Varshney, P. P. Das, P. D. Clough, S. Chattopadhyay, An effective low-dimensional software code representation using bert and elmo, in: 2022 IEEE 22nd International Conference on Software Quality, Reliability and Security (QRS), IEEE, 2022, pp. 763–774.
  - [34] L. Igual, S. Seguí, L. Igual, S. Seguí, Introduction to data science, Springer, 2017.