

A Qualitative Investigation into LLM-Generated Multilingual Code Comments and Automatic Evaluation Metrics

Katzy, Jonathan; Huang, Yongcheng; Panchu, Gopal Raj; Ziemlewski, Maksym; Loizides, Paris; Vermeulen, Sander; van Deursen, Arie; Izadi, Maliheh

DOI

[10.1145/3727582.3728683](https://doi.org/10.1145/3727582.3728683)

Licence

CC BY-NC-SA

Publication date

2025

Document Version

Final published version

Published in

PROMISE 2025 - Proceedings of the 2025 21st International Conference on Predictive Models and Data Analytics in Software Engineering

Citation (APA)

Katzy, J., Huang, Y., Panchu, G. R., Ziemlewski, M., Loizides, P., Vermeulen, S., van Deursen, A., & Izadi, M. (2025). A Qualitative Investigation into LLM-Generated Multilingual Code Comments and Automatic Evaluation Metrics. In *PROMISE 2025 - Proceedings of the 2025 21st International Conference on Predictive Models and Data Analytics in Software Engineering* (pp. 31-40). ACM.
<https://doi.org/10.1145/3727582.3728683>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



A Qualitative Investigation into LLM-Generated Multilingual Code Comments and Automatic Evaluation Metrics

Jonathan Katzy
Delft University of Technology
Delft, The Netherlands
J.B.Katzy@TUDelft.nl

Yongcheng Huang
Delft University of Technology
Delft, The Netherlands
Y.Huang-51@Student.TUDelft.nl

Gopal-Raj Panchu
Delft University of Technology
Delft, The Netherlands
G.G.S.Panchu@Student.TUDelft.nl

Maksym Ziemlewski
Delft University of Technology
Delft, The Netherlands
M.Ziemlewski@Student.TUDelft.nl

Paris Loizides
Delft University of Technology
Delft, The Netherlands
P.Loizides@Student.TUDelft.nl

Sander Vermeulen
Delft University of Technology
Delft, The Netherlands
S.R.Vermeulen@Student.TUDelft.nl

Arie van Deursen
Delft University of Technology
Delft, The Netherlands
Arie.vanDeursen@TUDelft.nl

Maliheh Izadi
Delft University of Technology
Delft, The Netherlands
M.Izadi@TUDelft.nl

Abstract

Large Language Models are essential coding assistants, yet their training is predominantly English-centric. In this study, we evaluate the performance of code language models in non-English contexts, identifying challenges in their adoption and integration into multilingual workflows. We conduct an open-coding study to analyze errors in code comments generated by five state-of-the-art code models, CodeGemma, CodeLlama, CodeQwen1.5, GraniteCode, and StarCoder2 across five natural languages: Chinese, Dutch, English, Greek, and Polish. Our study yields a dataset of 12,500 labeled generations, which we publicly release. We then assess the reliability of standard metrics in capturing comment *correctness* across languages and evaluate their trustworthiness as judgment criteria. Through our open-coding investigation, we identified a taxonomy of 26 distinct error categories in model-generated code comments. They highlight variations in language cohesion, informativeness, and syntax adherence across different natural languages. Our analysis shows that, while these models frequently produce partially correct comments, modern neural metrics fail to reliably differentiate meaningful completions from random noise. Notably, the significant score overlap between expert-rated correct and incorrect comments calls into question the effectiveness of these metrics in assessing generated comments.

CCS Concepts

• **Computing methodologies** → **Natural language generation.**

Keywords

Multilingual, Comment Generation, Open Coding, Large Language Models, Qualitative Evaluation



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

PROMISE '25, Trondheim, Norway

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1594-5/2025/06

<https://doi.org/10.1145/3727582.3728683>

ACM Reference Format:

Jonathan Katzy, Yongcheng Huang, Gopal-Raj Panchu, Maksym Ziemlewski, Paris Loizides, Sander Vermeulen, Arie van Deursen, and Maliheh Izadi. 2025. A Qualitative Investigation into LLM-Generated Multilingual Code Comments and Automatic Evaluation Metrics. In *21st International Conference on Predictive Models and Data Analytics in Software Engineering (PROMISE '25)*, June 26, 2025, Trondheim, Norway. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3727582.3728683>

1 Introduction

The adoption of Large Language Models (LLMs) for code has led to numerous benefits. When adopting LLMs in the software development process, studies reveal that LLMs boost developer productivity [60], and facilitate diverse software engineering tasks [13, 11]. Even in education, LLMs have shown to be beneficial when it comes to teaching programming [5], and assisting developers of various skill levels [30].

While code models can enhance both learning and development, most research and development in this area focus on *English* contexts. This emphasis limits the adoption of code models for other languages, excluding non-English speakers from the benefits of LLMs in education [25], access to learning in their preferred language [2, 33], and integration of code models into workflows that often involve substantial non-English code [10, 35]. Beyond adoption challenges, non-English code is sometimes considered a “data smell” to be avoided when training code models [51] leading to less and less representation.

By non-English code, we refer to the presence of a language other than English in either the identifiers, literals, or comments. This is often a mixture of English and another language. While non-English code is not commonly used in open source projects, close source proprietary code bases often contain other languages [35].

We aim to evaluate the performance of several state-of-the-art code LLMs in non-English contexts, assess their effectiveness and identify limitations that may impede their adoption in low-resource languages. More specifically, we evaluate five code LLMs—CodeGemma [42], CodeLlama [38], CodeQwen1.5 [44],

GraniteCode [28], and StarCoder2 [24]—across five natural languages: *English, Dutch, Chinese, Greek, and Polish*. To assess their performance in each language, we focus on comment generation, a task that serves as a proxy for code comprehension and linguistic adaptation. We take both *qualitative* and *quantitative* approaches to assess the strengths and weaknesses of the selected models.

For the qualitative investigation, we conduct an open-coding investigation into errors and accuracy of our target models. This gives us an overview of errors and insight into the ‘correctness’ of the comments. Gaining deeper insight into errors made in the generation is essential to understand what the current challenges are with regards to trust [54], education [31] and the willingness of practitioners to adopt models into their non-English workflows [10].

For the quantitative evaluation of comment generation, we investigate how well common metrics capture the performance of LLMs in non-English comment generation. We use the expert-labeled comments as a baseline to compare the metrics scores to and analyze their ability to differentiate between correct and incorrect predictions, as well as their ability to differentiate real generations from noise. More specifically, we answer the following questions;

- What are common errors when LLMs generate comments in different languages?
- How well do LLMs generate comments in non-English code?
- Do common evaluation metrics give a trustworthy evaluation of predictions?
- Do common evaluation metrics effectively differentiate between correct and incorrect predictions?

Our investigation resulted in a taxonomy of 26 distinct errors made by LLMs in 5 different languages. We show that linguistic errors increase dramatically when working in non-English contexts (up to 15.1×), as well as the likelihood that models produce incoherent generations. These increases are higher than the increases in semantic errors, which increased by 3.7× at most. This has strong implications for the adoption of code models into the workflow in non-English settings.

When investigating the metrics, we see that neural metrics struggle to differentiate real predictions from random noise. We showed that there is a substantial overlap between the scores assigned to random noise and to real generations when working with neural metrics. We also see that neural metrics give an overall higher score to all predictions, especially in Chinese, even when human annotators rated most predictions as incorrect.

2 Related Work

Multilingual LLMs in NLP. Studies on LLMs in NLP settings have shown that they perform poorly in non-English and resource-poor natural languages. The performance of several models investigated in Indonesian and African languages continues to fall short, producing responses of lower quality compared to other high-resource languages such as English [17, 32]. Further studies showed lower performance for non-English queries to chat models resulting in less accurate and lower quality responses compared to English queries, as well as issues arising with regard to cultural nuances when dealing with non-English queries [57]. Such findings are consistent with research suggesting a strong bias in multilingual settings when generating code based on non-English instructions. In the case of

Code LLaMA, a 37.8% drop in performance (Pass@1 metric) was observed when code generation was prompted in Chinese rather than English [53]. This performance gap emphasizes the underlying bias in LLMs trained predominantly on English-language datasets.

Comment/Documentation Generation. Automatic comment generation is a common task that has shown to help developers understand code [16]. In this task, language models are used to generate a natural language description of a piece of code. While it has shown to help developer productivity, most studies focus primarily on the performance of models in English code. Aside from evaluating models on comment generation, research has shown that for automated comment generation in general; practitioners expect automatically generated comments to be fluent and grammatically correct (82%), be concise (88%), and accurate (88%) [10] in order to include them in their workflows.

Qualitative Analyses of LLM Outputs. Previous works on translation, and specializations on automatic translations have led to the development of Taxonomies like SCATE [46] and evaluation frameworks such as MQM [26], used for natural language tasks. These provide a basis for error classification, however they do not capture all intricacies of code related tasks. Further qualitative studies of code generation errors have identified common bugs in LLM-generated code [41, 12] as well as given important insights into user preference for model invocations and acceptance rates [12], however, they make the assumption that code is always written in English, and they focus primarily on the correctness of a program, rather than specific errors in the language.

Non-English Code. Despite the global prevalence of code written in different languages, non-English code remains significantly under-researched. To the best of our knowledge, only one study has looked at the distribution of non-English language in code repositories [35]. They show that there is a significant increase in the use of German in both comments and identifier names when looking at industry projects, compared to open source projects, which showed to have almost no non-English text present. Datasets that are intended to be used for the training of LLMs, do mention the presence of other languages in the code, however, do not give any more specific information [24].

for code LLMs, most research focuses on English contexts; even researchers who are not native English speakers frame their studies within an English paradigm [23], leading to an under-representation of non-English scenarios. Reports from GitHub show that a significant number of non-English speakers contribute code to the platform each year, highlighting the importance of studying code generation in non-English contexts [8].

Metrics Reliability. The ability to evaluate the quality of text generated by LLMs is important to fairly evaluate and compare different models. There are, however, some challenge that arise when attempting to evaluate natural language generation programmatically. First, many early methods for evaluating textual generations come from the translation domain, where a generation can be compared to multiple reference texts in order to score them appropriately [34, 20, 3]. These methods, however, mainly look at an overlap of letters in the target generation with the reference generations. This creates issues when comparing the correctness of

the meaning behind the text, which has been proposed to be solved with embedding based metrics [58, 59] or model based metrics [56]. This has led to the development of frameworks that can be used to compare metrics to each other [55] as well as studies measuring the correlation between different metrics [6] and their resilience to simple perturbations [39]. These studies give an overview of what is to be expected from a metric to be a good judge of generated text. We summarize the requirements as follows: A metric must differentiate correct from incorrect samples, agree with human expert expectations, and be able to detect simple perturbations.

3 Non-English Characteristics

To gain an overview of the differences in the selected languages compared to English, we list the defining characteristics and how they may have an influence on the outputs of a model.

Morphosyntactic Complexity. Morphosyntactic features govern how words change form (morphology) and combine into sentences (syntax). Chinese’s uninflected monomorphemic system eliminates grammatical markings like tense/case [36], while Polish uses seven grammatical cases and gendered conjugations. Greek combines gendered nouns with flexible word order through inflectional endings, and Dutch employs gendered articles (de/het). These differences add a layer of complexity to the generation of languages compared to English. Furthermore, Polish and Greek are both flexible when it comes to following a Subject Verb Object (SVO) structure of a sentence, making the outputs context dependent [9, 49]

Orthographic Challenges. Orthographic systems determine how languages represent speech in writing. Greek’s technical symbols (e.g., α , β) and diacritics (tonos/diaeresis) coexist with Polish’s nine accented characters (ą, ó, etc.), while Chinese uses logograms instead of alphabetic writing. These add a layer of complexity to the tokenization of text, and require language-specific tokens. Furthermore, these languages face informal script variations: Greeklish (Latinized Greek), diacritic omission in Polish code comments, and Chinese character simplification. These are informal ways of writing a language, often adopted in online spaces which add complexity for the model to learn [47]. Finally, Greek has also been adopted as a common alphabet for use in mathematics, this adds another source of noise that could reduce a model’s ability to model Greek.

Semantic Nuances. Semantic features require understanding the meaning beyond direct translation. Chinese requires contextual interpretation of idioms [37], Dutch incorporates English loanwords through prolonged contact [50], while Polish/Greek use inflection-enabled word order flexibility to convey pragmatic emphasis. Such features demand cultural and syntactic awareness beyond direct translation equivalents, challenging LLMs’ ability to preserve intended meaning during code comment generation.

4 Approach

4.1 Data Collection

To gather a code dataset for each natural language, a list of the 2,500 most common words in each language was used¹, based on the OpenSubtitles2016 dataset [21]. For each word, the GitHub API

¹<https://github.com/oprogramador/most-common-words-by-language>

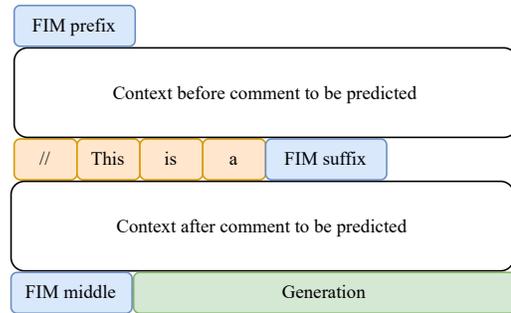


Figure 1: Example input used for inference

was queried to gather 100 files containing the specific word. The collected files were then de-duplicated.

Due to varying architectures, each model has its own maximum token length. To ensure a fair comparison between models, we filter the candidate files by their length after tokenization. We then select only files where all of the context fits within the context of the smallest model (CodeGemma [42], 4,128 tokens). To define the inference limit, we calculated the average comment length plus three standard deviations of the ground truth of comments. We filtered out all files where the context length + average length + 3 standard deviations was greater than 4,128 tokens to ensure that all files could be correctly predicted by all models.

For each file, the comments were extracted using regular expressions. After extracting the comments, the langdetect toolkit [52] was used to verify that the comments were written in the appropriate language. The target comments were further filtered manually to remove comments which were references to licenses, contained only the authors’ personal/contact information, TODOs, or were auto-generated. From the remaining files, comments with less than 10 tokens in all tokenizers were filtered out. Finally, 500 random files were sampled for each language, resulting in 2,500 generated comments for each language.

Comment Generation. To control the number of variables to account for, such as finetuning tasks and prompting template, we evaluate base models for their ability to generate comments. To do this we use the Fill In the Middle [4] (FIM) input format. This is the same task that is used to train the base LLMs, giving us the best idea of their core performance. In order to ensure that the models will generate predictions in the correct language, we add the first 3 tokens of the ground truth in order to prime the model for the correct language. Not doing so results in almost exclusively English comment generation. To give an overview of the input format we give an example in Figure 1.

4.2 Qualitative Analysis

Open Coding. To label the generations, an open coding methodology was applied. An initial set of errors based on previous work [12, 46, 26] was selected and iteratively improved in multiple rounds of discussion between 6 authors, for each language there was at least one author that had a native speaker level of understanding of the language to ensure that linguistic nuances, grammar correctness, and semantic accuracy were accurately judged.

An initial subset of 200 files per language was selected for analysis, similar to the final dataset. The generated comments were manually compared to the originals to classify any possible errors according to the initial hypothetical error taxonomy. The results of this labeling iteration were discussed in meetings between 6 authors, resulting in adjustments to the error taxonomy. After each iteration, 6 authors discussed their findings regarding the error taxonomy and adapted the taxonomy to reflect any changes that needed to be made according to newly identified categories. In the process of adapting the taxonomy, inclusion and exclusion criteria were defined for each category to clarify when an error should be labeled within that category. This evaluation process was repeated with new data until no further adjustments to the error taxonomy were suggested, resulting in five iterations until the error taxonomy no longer changed. The process of generating a final taxonomy, and labeling all files cost 500 person-hours.

The resulting error taxonomy was then used for error labeling on our final dataset. This dataset contains 500 files per language for a total of 2,500 files, with one inference from each model resulting in 5 inferences per file. This totals 12,500 labeled comments.

Expert Accuracy. As well as recording the errors made by a model, we also use an evaluation of a native speaker with programming experience to determine whether a prediction is correct, partially correct, or completely incorrect. In practice, model suggestions for code completion may be accepted, but will be edited afterward to better fit the problem at hand [29, 19, 12]. Including partially correct as a label gives us a more nuanced approach for situations where the models are close to a valid answer but still make minor mistakes. This accounts for developers accepting partially correct completions that they later edit themselves [12].

We define the categories of correct, partially correct, and incorrect as follows. A correct prediction must include all of the information in the original comment but may also include additional, relevant, information that was omitted. A partially correct comment is a comment where the content is largely correct but there is a small error, such as an incorrect variable name, or spelling mistake. All other comments were considered to be incorrect.

4.3 Quantitative Analysis

For the numerical analysis of the comments, and validation of the metrics, we treat the problem as a text generation problem. We have a prediction given by the model and a reference comment we compare it to. For this evaluation we use a combination of word level and neural metrics. To validate the metrics' ability to score a prediction, we focus on the requirement of a metric to reduce the score of a generation significantly due to a perturbation [39]. We push these perturbations to the extreme. We generate two types of random noise. First we generate "uniform noise" which is any token from the models' tokenizer, equal in length to the target comment in tokens, and "targeted noise" which are random tokens sampled from the surrounding context. We give an example of the noise and scoring of a generation in Figure 2.

Word Level Metrics. For the word-level metrics, we use the standard metrics used in code-related studies: BLEU [34], ROUGE [20], and METEOR [3]. BLEU measures precision by comparing n-gram

matches between the generated and reference summaries while penalizing overgeneration with a brevity penalty. ROUGE focuses on recall by measuring the overlap of n-grams between the reference and the generated texts. ROUGE-1 considers unigrams, ROUGE-2 considers bigrams, and ROUGE-L emphasizes the longest common subsequence to assess fluency and logical consistency. For this investigation, we focus on ROUGE-L. METEOR combines precision, recall, and synonym-based matching to provide the score.

Embedding Based Metrics. These metrics calculate the cosine similarity between embeddings. BERTScore [58] uses a BERT [15] model to generate the embeddings. The BERT configuration varies according to the target language. For English tasks, a 24-layer RoBERTa large model is used, while for Chinese tasks, a 12-layer BERT Chinese model is utilized. For other languages, it relies on the 12-layer cased multilingual BERT (BERT_{multi}) model. BERT_{multi} was trained on 104 languages using the Wikipedia corpora, including Greek, Polish, and Dutch. This model, however, has not been trained on code. CodeBERTScore [59] is similar to BERTScore, however, it was specifically developed for code evaluation, it differs from BERTScore, by using the CodeBERT model [7] and ignoring some keywords and syntax tokens common throughout code.

Model Based Metrics. BARTScore [56] employs BART [18], a transformer model that evaluates the quality of generated text through likelihood scores assigned by the model. The BART models is trained by both corrupting the input, and introducing arbitrary noise into the input text, after which it reconstructs the original input. While specific details regarding training data are not explicitly provided, some indirect information suggests that it includes datasets such as books and Wikipedia entries—similar to those used in RoBERTa's pre-training process [22]. Similar to BERTScore, code is not specifically mentioned in the training procedure.

5 Models

We focus our investigation on open-weight models. This gives us full control over the generations. Ensuring that no other pre-processing steps are performed unknown to us, and allows for reproducible research, which is not guaranteed when using proprietary solutions [40]. More specifically, we focus on models within the 7B-8B parameter range, the largest amount of parameters that have releases of all 5 model families. To understand the prevalence of our target languages in the training procedure of these models, we present a brief overview.

CodeQwen1.5-7B. CodeQwen1.5-7B extends Qwen1.5 with an English and Chinese dataset, supplemented by 3T tokens of code data spanning 92 programming languages, though the distribution of natural languages in this corpus is unspecified [44, 45].

StarCoder2-7B. StarCoder2-7B is trained on The Stack V2 dataset², totaling 525.5B tokens of programming languages and documentation, with additional natural language data from Wikipedia, Stack-Overflow, Arxiv, Free Programming Books³, and OpenWebMath, covering all languages in our study [24].

²<https://huggingface.co/datasets/bigcode/the-stack-v2>

³<https://github.com/EbookFoundation/free-programming-books>

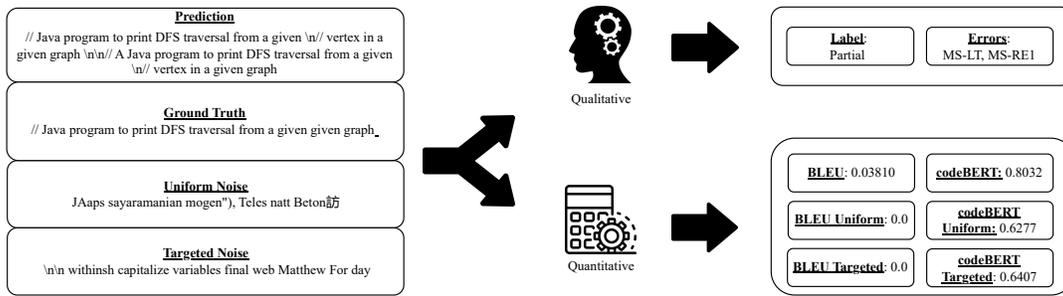


Figure 2: Example Evaluation of a Prediction

Granite-8B-code-base. Granite-8B-code-base is trained on a collection of programming languages datasets, including GitHub Code Clean⁴, StarCoderData⁵, and GitHub issues, combined with natural language datasets from StackExchange, Arxiv, and OpenWebMath. English is predominant, with non-English data removed [28].

CodeLlama-7B. CodeLlama-7B builds on Llama2, whose training corpus is 89.7% English, with minor representation of Chinese, Dutch, and Polish. CodeLlama is further trained on publicly available code, incorporating 8% natural language data related to code and 7% from natural language dataset, though specific origins remain undisclosed [38, 48].

CodeGemma-7B. CodeGemma-7B originates from the Gemma family, trained on 6T tokens inspired by Gemini, which training details remain undisclosed. It is fine-tuned on a dataset consisting of 80% code and 20% natural language, predominantly English, with no specified multi-lingual focus. However, dataset sources are also undisclosed [42, 43].

6 Results

6.1 Error Taxonomy

From the open coding investigation, we arrive at 4 main categories of error in comment generation. We present these categories and all their subcategories in Table 1. We use this taxonomy to answer RQ1, What are common errors when LLMs generate comments in different languages? Semantic errors, the most common category, refer to errors that relate to the meaning of the comment. This can relate to the amount of detail provided to the model, whether the model is hallucinating, and whether code is included in/after the comment. Linguistic errors refer to errors made in the language itself. This can refer to grammatical errors, but also to which synonyms the model uses and whether it responds in the correct language. Model-specific errors are errors due to the behavior of the LLMs; this focuses on memorization, copying information from the context, and knowing when to stop.

Each of the four categories can be divided into subcategories. The most common in the Model-Specific category is the Late Termination (MS-LT) error, which occurs when the model continues to produce additional content after the relevant part of the comment has been completed. On the other hand, there is Early Termination

(MS-ET), which means that the prediction stops when the comment is not complete. Other error types are Copy Context (MS-CC) when the model copies the surrounding context verbatim, and Repetition (MS-RE), when the model repeats what it has already generated, this can happen in two ways. The model can repeat a pattern, but make minor changes such as increasing a number (MS-RE1) or the model can repeat a token or set of tokens exactly the same (MS-RE2). Incoherent Generation (MS-IG) occurs when the generation consists of random words or symbols with no logic between them. Memorization (MS-ME) is assigned when the model returns personally identifiable information (MS-ME1), URL (MS-ME2), or the prediction is the same as the erroneous ground truth (MS-ME3), indicating that the comment was in the training set (e.g. the original comment contained a grammatical error and a model repeated that, or the original comment did not refer to the code correctly and the model predicted exactly the same).

Linguistic errors are divided into three subcategories. The first one is Grammar (LG-GR), where all grammatical errors are assigned. The second one, Incorrect Synonym (LG-IS), occurs when a model uses a similar word with an incorrect meaning in context. The last category is the Wrong Language (LG-WL) category. Due to the use of English loan words in many languages, the models will attempt to translate them, when they should not (LG-WL1). Similarly, the model can also generate most of, or a significant part of a comment in a language other than the target language (LG-WL2).

In the Semantic group, we distinguish several errors, such as the prediction with Missing Details (SE-MD) or being Too Specific (SE-TS). Many predictions also include code, either executable (SE-CS2) or commented out (SE-CS1). Subcategory for Hallucination generations is divided into three types of errors. Misplaced Facts (SE-HA1) occurs, when random facts that do not align with the expected content are present. Out of Context (SE-HA2), which is a hallucination not grounded in the provided context. and In Context (SE-HA3) Hallucinations, which are educated guesses based on the provided context. Finally, we have the Syntax errors (ST), which are all errors related to the syntax of the comments.

This taxonomy shows that the main issue with comments generated by LLMs is Semantic Errors, often in the form of Hallucinations (SE-HA). These errors account for the largest increase in errors for non-English languages compared to English. Furthermore, we see that the level of detail (SE-MD and SE-TS) is a cause of errors mainly in non-English generations. Looking at the Model Specific errors (MS) we see that for errors related to Memorization (MS-ME), the

⁴<https://huggingface.co/datasets/codeparrot/github-code-clean>

⁵<https://huggingface.co/datasets/bigcode/starcoderdata>

Table 1: Taxonomy of common errors found during the open coding process, all languages are listed separately, and summed up in the total. Sub-categories are accumulated in the count of the parent category.

Failure category with label ID	Total	Chinese	Dutch	English	Greek	Polish
MS Model-specific Errors	5,007	1,191	1,072	752	952	1,040
└ MS-IG Incoherent Generation	259	36	42	2	167	12
└ MS-CC Copy context	991	81	291	70	307	242
└ MS-ME Memorization	422	224	36	74	41	47
└└ MS-ME1 PII	236	71	32	68	34	31
└└ MS-ME2 URL	114	99	1	5	3	6
└└ MS-ME3 Training Data	72	54	3	1	4	10
└ MS-ET Early Termination	164	34	47	13	51	19
└ MS-LT Late Termination	2,324	660	405	491	296	472
└ MS-RE Repetition	847	156	251	102	90	248
└└ MS-RE1 Pattern Repetition	317	55	109	79	6	68
└└ MS-RE2 Verbatim Repetition	530	101	142	23	84	180
LG Linguistic Error	1,728	17	224	66	998	423
└ LG-GR Grammar	1,420	0	118	65	887	350
└└ LG-GR1 Plurality	17	0	1	0	15	1
└└ LG-GR2 Conjugation	169	0	8	0	59	102
└└ LG-GR3 Gender	130	0	23	0	84	23
└└ LG-GR4 Language Syntax	227	0	25	42	39	121
└└ LG-GR5 Cohesion	876	0	60	23	690	103
└ LG-IS Incorrect synonym	128	0	0	0	96	32
└ LG-WL Wrong language	180	17	106	1	15	41
└└ LG-WL1 Undesired translations	40	1	9	0	12	18
└└ LG-WL2 Incorrect language	140	16	97	1	3	23
SE Semantic error	8,333	2,915	1,350	786	1,759	1,523
└ SE-MD Missing Details	837	413	107	19	96	202
└ SE-TS Too Specific	168	91	21	1	13	42
└ SE-HA Hallucination	4,239	1,639	636	406	928	630
└└ SE-HA1 Misplaced Facts	325	79	105	42	8	91
└└ SE-HA2 Out of Context	550	385	80	12	57	16
└└ SE-HA3 In context	3,364	1,175	451	352	863	523
└ SE-CS Code Inclusion	3,009	718	574	355	722	640
└└ SE-CS1 Commented code	171	6	68	12	15	70
└└ SE-CS2 Runnable code	2,838	712	506	343	707	570
└ SE-OI Omitted Identifier	80	54	12	5	0	9
ST Syntax	84	31	2	21	5	25
└ ST-IF Incorrect comment format	84	31	2	21	5	25

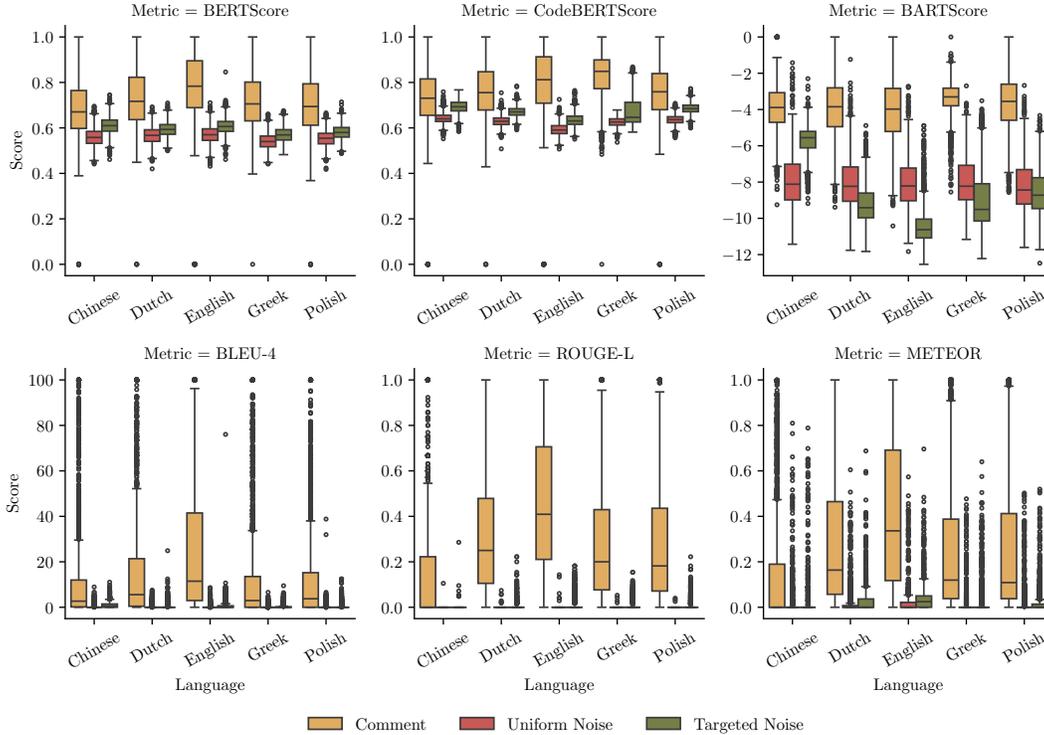
models are most likely to make them in English compared to other Western languages. Interestingly, Chinese has the highest rate of memorization for all categories. For language-specific errors (LG), we see that Chinese has no grammar errors, as discussed in Section 3. Furthermore, we notice that Dutch has the highest likelihood of returning predictions in the wrong language. We believe this is due to the similarity of Dutch to English, and the high amount of overlap between the two dictionaries. We also observe a high rate of incorrect synonyms used in Polish compared to all other languages and a high rate of language syntax errors for English,

compared to all other language errors made by the model. Finally, in the Syntax errors (ST) we see that the models' ability to adhere to comment syntax is dependent on the language it is using.

We further compare languages on a per-language basis and see that Greek has the highest number of errors for the language grammar, as well as the use of incorrect synonyms. This amounts to an increase of $15.1\times$ the error rate compared to English. This also appears in the MS category, where Greek has the highest tendency to generate incoherent outputs. On the other hand, Greek is the

Table 2: Expert accuracy for all models and languages (✓ correct, ~ partial, ✗ incorrect)

	CodeGemma			CodeLlama			CodeQwen 1.5			GraniteCode			StarCoder2		
	✓	~	✗	✓	~	✗	✓	~	✗	✓	~	✗	✓	~	✗
Chinese	55	221	224	63	201	236	84	200	216	67	175	258	20	218	262
Dutch	298	98	104	246	95	159	158	109	233	253	80	167	215	107	178
English	406	76	18	266	162	72	268	183	49	339	123	38	320	145	35
Greek	259	117	124	118	119	263	78	165	257	132	130	238	118	131	251
Polish	271	146	83	176	164	160	159	157	184	205	151	144	167	154	179

**Figure 3: Metric scores comparing LLM generated comments, to random samples of tokens using two separate distributions.**

language that is least likely to generate repetitions out of all analyzed languages. We also see that Dutch, Greek, and Polish are more likely to copy the surrounding context into the comment, while Chinese does it at a rate similar to English.

6.2 Expert Accuracy

To understand how well the chosen models perform at generating code comments in languages they were not intended to be used for (RQ2), the experts labeled each generation as being either “Correct”, “Partially Correct”, or “Incorrect”. We give an overview of the scores in Table 2. Here we see that while models claim to be only intended to be used for English, and sometimes Chinese, we see that they generate at least partially accurate completions in more than half of the cases for Dutch, Greek, and Polish. Interestingly, we see that although Chinese is the second most common language present in training corpora, it is the worst-performing language overall. On a model level, we see that most models follow the same trend, they

perform best in English, then Dutch or Polish, followed by Greek, and finally Chinese. However, for CodeGemma, we see that the performance of Greek is doubled in comparison to other models.

6.3 Scoring Collections

We answer RQ3 by analyzing the scores the selected metrics assign to both predictions and the random noise we generated. An overview is given in Figure 3. Here we look at two factors identified as belonging to a good metric, alignment with expert evaluations, and the ability to differentiate between noise and a real prediction.

Looking at only the English results, we see that the scores assigned by the embedding based metrics align closely with the expert evaluation (around 0.8). However, when looking at other languages, the embedding-based metrics score the predictions slightly lower, even when the expert evaluation scores the generations much lower (especially for Chinese). The word-based metrics on the other hand

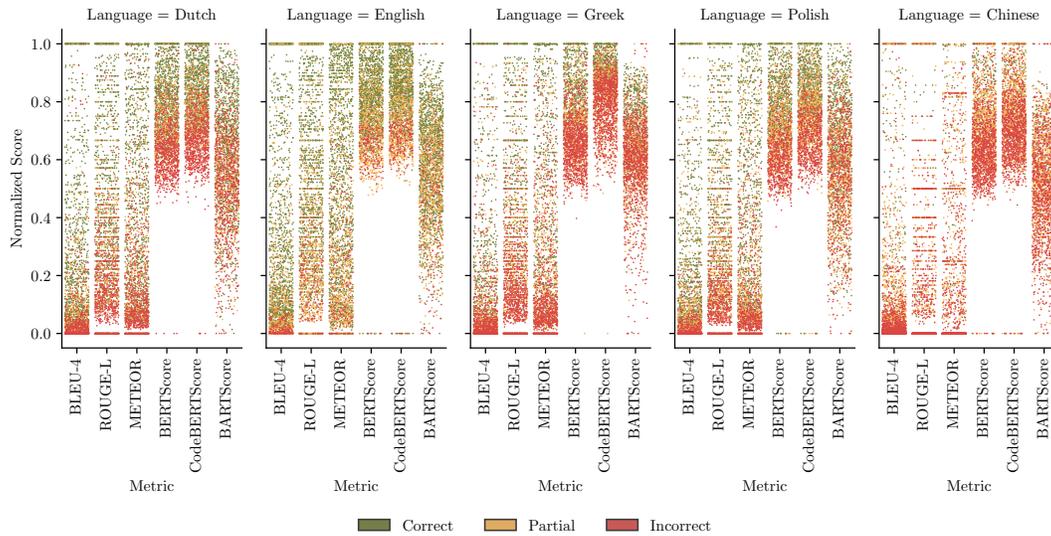


Figure 4: Strip plot, showing the scores assigned to comment generations by different metrics.

score predictions lower in general, however, do match the differences in performance observed in the expert accuracy.

Finally, when looking at noise, we see that for all neural metrics, there is a significant overlap between the scores assigned to noise and the scores assigned to real generations. The BARTScores have a lower overlap compared to the embedding based metrics, however, interestingly, the targeted noise performed worse than the uniform noise for all languages except Chinese. Finally, we see a general trend of CodeBERTScore scoring predictions higher than BERTScore, hinting that the training procedure is a more important factor in a higher BERTScore than generations being studied. For classic metrics, there is a far lower overlap between noise and real comments, and most noise is scored close to 0 with the METEOR score showing the largest number of outliers towards high scores.

6.4 Differentiating Correct From Incorrect

To answer RQ4, Do common evaluation metrics effectively differentiate between correct and incorrect predictions? We plot the scores of selected metrics in Figure 4. Here we show a strip plot of the score for each generated comment and color it by the expert evaluation “Correct” (green), “Partial” (orange), and “Incorrect” (red). Each point in a stripe represents a single generation. To demonstrate the spreads of different scores, we normalize all scores between 0 and 1. This shows us if there is a clear separation between the three colors at a glance. We see a difference in the distribution between neural models, and classical models. Classical models have the majority of their data points scored low and have a tail of scores that score high. On the other hand, neural metrics, especially the embedding-based metrics, have their scores concentrated towards the top, with few points in the lower half of the plot. Looking at all the strip plots, we see no clear separation between correct, partially correct, or incorrect predictions. While there is a slight trend to score correct predictions higher, a large overlap between the three colors persists, showing that the metrics do not separate them accurately.

7 Discussion

The results show two main trends. First, we see that the performance of non-English comment generation is often acceptable. However, the generated errors have strong implications for their use by practitioners. Secondly, using neural metrics for automatic evaluation gives an inflated view of an LLM’s performance.

We show that the number of errors relating to factuality of the comments increased by at most a factor of 3 compared to English, the number of errors related to language increased by a factor of 15.1. This larger decrease in the models’ ability to generate accurate natural languages damages the trust users have in the model, and is one of the main requirements practitioners name for automatic comment generation [10]. Furthermore, the level of detail (SE-MD) drops significantly (up to 21.8× for Chinese) which is another factor limiting the adoption of models by practitioners.

Looking at the results of the metric evaluations, we see a worrying trend for neural metrics. The inability of a metric to differentiate random noise from a real generation makes the other evaluations they give questionable. Comparing the expert accuracy scores with the results of the metrics we see that the human evaluators rated English comments correct 5.5× more often than Chinese comments. However, for the neural metrics, the Chinese score is marginally lower than the English score. More worryingly, For CodeBERTScore and BARTScore, the models score Greek better than all other languages, despite being second to last for human raters. Looking at the scores calculated by traditional metrics, we see that while the scores are lower overall, the respective ranking between the different languages matches more closely with the expert evaluators.

Our findings highlight several recommendations for improving multilingual code model development. First, model developers should prioritize diversifying training data to include more non-English code, particularly in languages like Chinese where performance degradation is most severe. Second, evaluation frameworks should require manual checks, as the results demonstrate

that neural metrics fail to reliably capture human judgments of comment quality across languages. We recommend maintaining human evaluators in assessment pipelines, particularly when evaluating non-English outputs, until more reliable automated metrics can be developed. Practitioners working with these models should be careful and conduct verification steps when generating comments or documentation in non-English languages. In addition, there is a need to quantify the distribution of non-English code in real-world development scenarios to gain more insight into the impacts of the problem of non-English code generation, and to assess the need to align model training for practical environments.

8 Threats to Validity

Memorization Risk. Most of the models under investigation do not publish their training data. Therefore, it cannot be known whether memorization took place and influenced the answers of the models for the given tasks. We are unable to control for this due to the prevalence of non-disclosed training datasets. In situations where the memorization was very obvious, we marked it as an error category. We implore open-source models to be released with their training data so research can be conducted fairly.

Labeling Bias. This study ran the risk of labeling bias due to a limited amount of experts labeling errors per language. We mitigate this bias by defining clear inclusion criteria per error label and having iterative discussions regarding these classifications between six authors. Labeling bias remains a risk in a multilingual context despite our mitigation efforts; therefore, we release all labeled data for validation to the public.

Limited Model Size. This investigation focused on models in the 7B-8B parameter range. While larger models tend to give better outputs, we limited the model size to the largest size available across all architectures under investigation. This prevents us from including confounding factors; which skew the results of the investigation.

9 Future Work

Our study evaluates the performance of five popular models for comment generation in multiple languages and highlights opportunities for future research.

Primarily, we observed a gap in the investigation of the multilingual nature of code. Although it is recognized within the community that code can be written in non-English languages, the extent to which this occurs on a large scale remains largely unknown. Additionally, there is limited research concerning the distribution of non-English language elements within code; these may appear in comments, function names, variable names, or strings, and often coincide with English to varying degrees. This complexity renders the detection of such languages a challenging problem. We propose future research focuses on mapping and annotating other languages present in extensive training datasets for code LLMs.

In our investigation, we also found that neural methods tend to struggle when distinguishing between correct and incorrect predictions. Furthermore, comparisons based solely on average values can be misleading when compared with random noise. These results closely resemble other works on multilingual embeddings of code tokens [14]. We believe that changing the training objectives

to include negative sampling, comparable to word2vec [27] is a promising avenue for future work.

Recent research on employing LLMs to replace manual annotation of software engineering artifacts [1] has shown that while the annotation generated by LLMs can align with human annotations, they are costly to be evaluated. To aid in the development and evaluation of these systems, we invite the community to use our publicly available dataset to compare the model outputs with human experts' annotations.

10 Conclusion

We explored the ability of LLMs to generate comments in Chinese, Dutch, English, Greek, and Polish. We performed open-coding to create a taxonomy of model failures when generating comments. Expert evaluators assessed the correctness of these comments and analyzed the effectiveness of various metrics in distinguishing noise from real comments, as well as correct from incorrect ones.

Our qualitative analysis revealed notable behaviors across languages, highlighting which areas are most important to address when designing systems for non-English contexts. Our metric analysis found shortcomings in neural metrics: cosine similarity-based metrics struggled to differentiate noise from genuine predictions, and all neural metrics failed to reliably separate correct from incorrect comments. To support further research, we release our 12,500 labeled predictions including five languages and five models⁶.

References

- [1] Toufique Ahmed, Premkumar Devanbu, Christoph Treude, and Michael Pradel. 2025. Can llms replace manual annotation of software engineering artifacts? In *2025 IEEE/ACM 22nd International Conference on Mining Software Repositories*.
- [2] John Airey and Cedric Linder. 2006. Language and the experience of learning university physics in sweden. *EUROPEAN JOURNAL OF PHYSICS Eur. J. Phys.*, 27, (May 2006), 553–560. doi:10.1088/0143-0807/27/3/009.
- [3] Satanjeev Banerjee and Alon Lavie. 2005. Meteor: an automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 65–72.
- [4] Mohammad Bavarian, Heewoo Jun, Nikolas Tezak, John Schulman, Christine McLeavey, Jerry Tworek, and Mark Chen. 2022. Efficient training of language models to fill in the middle. *arXiv preprint arXiv:2207.14255*.
- [5] Samuel Boguslawski, Rowan Deer, and Mark G Dawson. 2025. Programming education and learner motivation in the age of generative ai: student and educator perspectives. *Information and Learning Sciences*, 126, 1/2, 91–109.
- [6] A. R. Fabbri, Wojciech Kryscinski, Bryan McCann, Richard Socher, and Dragomir R. Radev. 2020. Summeval: re-evaluating summarization evaluation. *Transactions of the Association for Computational Linguistics*, 9, 391–409.
- [7] Zhangyin Feng et al. 2020. CodeBERT: a pre-trained model for programming and natural languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Trevor Cohn, Yulan He, and Yang Liu, (Eds.) Association for Computational Linguistics, Online, (Nov. 2020).
- [8] GitHub. 2022. The state of the octoverse: global tech talent report 2022. Accessed: 2024-11-13. (2022). octoverse.github.com/2022/global-tech-talent.
- [9] Hubert Haider and Luka Szucsich. 2022. *Theoretical Linguistics*, 48, 1-2, 1–39. doi:doi:10.1515/tl-2022-2035.
- [10] Xing Hu, Xin Xia, David Lo, Zhiyuan Wan, Qiuyuan Chen, and Thomas Zimmermann. 2022. Practitioners' expectations on automated code comment generation. In *Proceedings of the 44th international conference on software engineering*.
- [11] Maliheh Izadi and et al. 2022. Codefill: multi-token code completion by jointly learning from structure and naming sequences. In *Proceedings of the 44th International Conference on Software Engineering*, 401–412.
- [12] Maliheh Izadi, Jonathan Katzy, Tim Van Dam, Marc Otten, Razvan Mihai Popescu, and Arie Van Deursen. 2024. Language models for code completion: a practical evaluation. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, 1–13.

⁶<https://huggingface.co/datasets/AISE-TUDELFT/multilingual-code-comments>

- [13] Juyong Jiang, Fan Wang, Jiashi Shen, Sungju Kim, and Sunghun Kim. 2024. A survey on large language models for code generation. *arXiv preprint arXiv:2406.00515*.
- [14] Jonathan Katzy, Maliheh Izadi, and Arie van Deursen. 2023. On the impact of language selection for training and evaluating programming language models. In *2023 IEEE 23rd International Working Conference on Source Code Analysis and Manipulation (SCAM)*. IEEE, 271–276.
- [15] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. Bert: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of naacl-HLT*. Vol. 1. Minneapolis, Minnesota, 2.
- [16] Junaed Younus Khan and Gias Uddin. 2022. Automatic code documentation generation using gpt-3. (2022). eprint: 2209.02235.
- [17] Fajri Koto, Nurul Aisyah, Haonan Li, and Timothy Baldwin. 2023. Large language models only pass primary school exams in Indonesia: a comprehensive test on IndoMMLU. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Houda Bouamor, Juan Pino, and Kalika Bali, (Eds.) Association for Computational Linguistics, Singapore, (Dec. 2023), 12359–12374. doi:10.18653/v1/2023.emnlp-main.760.
- [18] M Lewis. 2019. Bart: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension.
- [19] Jenny T. Liang, Chenyang Yang, and Brad A. Myers. 2024. A large-scale survey on the usability of ai programming assistants: successes and challenges. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering (ICSE '24)* Article 52. Association for Computing Machinery, Lisbon, Portugal, 13 pages. ISBN: 9798400702174. doi:10.1145/3597503.3608128.
- [20] Chin-Yew Lin. 2004. ROUGE: a package for automatic evaluation of summaries. In *Text Summarization Branches Out*. Association for Computational Linguistics, Barcelona, Spain, (July 2004), 74–81. <https://aclanthology.org/W04-1013>.
- [21] Pierre Lison and Jörg Tiedemann. 2016. OpenSubtitles2016: extracting large parallel corpora from movie and TV subtitles. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*. Nicoletta Calzolari et al., (Eds.) European Language Resources Association (ELRA), Portorož, Slovenia, (May 2016), 923–929. <https://aclanthology.org/L16-1147>.
- [22] Yinhan Liu. 2019. Roberta: a robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 364.
- [23] Irene López-Navarro, Ana Ibáñez Moreno, Miguel Angel Quintanilla, and Jesús Rey-Rocha. 2015. Why do i publish research articles in english instead of my own language? differences in spanish researchers' motivations across scientific domains. *Scientometrics*, 103, 939–976.
- [24] Anton Lozhkov et al. 2024. Starcoder 2 and the stack v2: the next generation. (2024). <https://arxiv.org/abs/2402.19173> arXiv: 2402.19173 [cs. SE].
- [25] Wenhan Lyu, Yimeng Wang, Tingting (Rachel) Chung, Yifan Sun, and Yixuan Zhang. 2024. Evaluating the effectiveness of llms in introductory computer science education: a semester-long field study. In *Proceedings of the Eleventh ACM Conference on Learning @ Scale (L@S '24)*. Association for Computing Machinery, Atlanta, GA, USA, 63–74. ISBN: 9798400706332.
- [26] Valerie R Mariana. 2014. *The Multidimensional Quality Metric (MQM) framework: A new framework for translation quality assessment*.
- [27] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. (2013). arXiv: 1301.3781.
- [28] Mayank Mishra et al. 2024. Granite code models: a family of open foundation models for code intelligence. (2024). arXiv: 2405.04324 [cs. AI].
- [29] Hussein Mozannar, Gagan Bansal, Adam Fourney, and Eric Horvitz. 2024. Reading between the lines: modeling user behavior and costs in ai-assisted programming. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)* Article 142. Association for Computing Machinery, Honolulu, HI, USA, 16 pages. ISBN: 9798400703300. doi:10.1145/3613904.3641936.
- [30] Daye Nam, Andrew Macvean, Vincent Hellendoorn, Bogdan Vasilescu, and Brad Myers. 2024. Using an llm to help with code understanding. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, 1–13.
- [31] Rajathurai Nishanthi. 2020. Understanding of the importance of mother tongue learning. *International Journal of Trend in Scientific Research and Development*.
- [32] Jessica Ojo, Kelechi Ogueji, Pontus Stenertorp, and David Ifeoluwa Adelani. 2024. How good are large language models on african languages? (2024).
- [33] Yogendra Pal and Sridhar Iyer. 2015. Classroom versus screencast for native language learners: effect of medium of instruction on knowledge of programming. In *Proceedings of the 2015 ACM Conference on Innovation and Technology in Computer Science Education (ITICSE '15)*. Association for Computing Machinery, Vilnius, Lithuania, 290–295. ISBN: 9781450334402. doi:10.1145/2729094.2742618.
- [34] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 311–318.
- [35] Timo Pawelka and Elmar Juergens. 2015. Is this code written in english? a study of the natural language of comments and identifiers in practice. In *2015 IEEE International Conference on Software Maintenance and Evolution (ICSME)*.
- [36] Edwin George Pulleyblank. 1995. *Outline of classical Chinese grammar*.
- [37] Jipeng Qiang, Yang Li, Chaowei Zhang, Yun Li, Yi Zhu, Yunhao Yuan, and Xindong Wu. 2023. Chinese idiom paraphrasing. *Transactions of the Association for Computational Linguistics*, 11, (July 2023), 740–754.
- [38] Baptiste Roziere et al. 2023. Code llama: open foundation models for code. *arXiv preprint arXiv:2308.12950*.
- [39] Ananya B. Sai, Tanay Dixit, Dev Yashpal Sheth, Sreyas Mohan, and Mitesh M. Khapra. 2021. Perturbation CheckLists for evaluating NLG evaluation metrics. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, (Eds.) Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, (Nov. 2021), 7219–7234.
- [40] Harald Semmelrock, Tony Ross-Hellauer, Simone Kopeinik, Dieter Theiler, Armin Haberl, Stefan Thalmann, and Dominik Kowald. 2024. Reproducibility in machine learning-based research: overview, barriers and drivers. *CoRR*, abs/2406.14325. <https://doi.org/10.48550/arXiv.2406.14325>.
- [41] Florian Tambon, Arghavan Moradi-Dakhel, Amin Nikanjam, Foutse Khomh, Michel C Desmarais, and Giuliano Antoniol. 2025. Bugs in large language models generated code: an empirical study. *Empirical Software Engineering*.
- [42] CodeGemma Team et al. 2024. Codegemma: open code models based on gemma. (2024). <https://arxiv.org/abs/2406.11409> arXiv: 2406.11409 [cs. CL].
- [43] Gemma Team et al. 2024. Gemma: open models based on gemini research and technology. (2024). arXiv: 2403.08295.
- [44] Qwen Team. 2024. Code with codeqwen1.5. (Apr. 2024). <https://qwenlm.github.io/blog/codeqwen1.5/>.
- [45] Qwen Team. 2024. Introducing qwen1.5. (Feb. 2024). <https://qwenlm.github.io/blog/qwen1.5/>.
- [46] Tezcan, Arda and Hoste, Veronique and Macken, Lieve. 2017. SCATE taxonomy and corpus of machine translation errors. eng. In *Trends in E-tools and resources for translators and interpreters*. Approaches to Translation Studies. Vol. 45. Pastor, Gloria Corpus and Durán-Muñoz, Isabel, (Ed.) Brill | Rodopi, 219–244.
- [47] Anastasios Toumazatos, John Pavlopoulos, Ion Androutsopoulos, and Stavros Vassos. 2024. Still all Greeklish to me: Greeklish to Greek transliteration. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*. Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue, (Eds.) ELRA and ICCL, Torino, Italia, (May 2024), 15309–15319.
- [48] Hugo Touvron et al. 2023. Llama 2: open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- [49] Dimitra Irimi Tzanidaki. 1995. Greek word order: towards a new approach. *UCL Working Paper in Linguistics*, 7, 247–277.
- [50] Lieke Verheijen and Roeland van Hout. 2022. Manifold code-mixing in computer-mediated communication: the use of english in dutch youths' informal online writing. *Ampersand*, 9, 100091. doi:<https://doi.org/10.1016/j.amper.2022.100091>.
- [51] Antonio Vitale, Rocco Oliveto, and Simone Scalabrino. 2024. A catalog of data smells for coding tasks. *ACM TOSEM*.
- [52] Tamás Kovács Florian Atzenhofer-Baumgartner Sandy Aoun Anguelos Nicolaou Daniel Luger Franziska Decker Florian Lamming Georg Vogeler. 2022. Langdetect (revision 0215f72). (2022). doi:10.57967/hf/0135.
- [53] Chaozheng Wang, Zongjie Li, Cuiyun Gao, Wenxuan Wang, Ting Peng, Hailiang Huang, Yuetang Deng, Shuai Wang, and Michael R. Lyu. 2024. Exploring multi-lingual bias of large code models in code generation. (2024).
- [54] Harry J. Witchel, Christopher I. Jones, Georgina A. Thompson, Carina E. I. Westling, Juan Romero, Alessia Nicotra, Bruno Maag, and Hugo D. Critchley. 2022. Spelling errors in brief computer-mediated texts implicitly lead to linearly additive penalties in trustworthiness. *Frontiers in Psychology*, 13, 873844.
- [55] Ziang Xiao, Susu Zhang, Vivian Lai, and Q Vera Liao. 2023. Evaluating evaluation metrics: a framework for analyzing nlg evaluation metrics using measurement theory. In *2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023*. Association for Computational Linguistics (ACL).
- [56] Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. Bartscore: evaluating generated text as text generation. *NeurIPS*, 34, 27263–27277.
- [57] Xiang Zhang, Senyu Li, Bradley Hauer, Ning Shi, and Grzegorz Kondrak. 2023. Don't trust chatgpt when your question is not in english: a study of multilingual abilities and types of llms. (2023). arXiv: 2305.16339 [cs. CL].
- [58] Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: evaluating text generation with bert. In *International Conference on Learning Representations*.
- [59] Shuyan Zhou, Uri Alon, Sumit Agarwal, and Graham Neubig. 2023. CodeBERTScore: evaluating code generation with pretrained models of code. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Houda Bouamor, Juan Pino, and Kalika Bali, (Eds.) Association for Computational Linguistics, Singapore, (Dec. 2023), 13921–13937.
- [60] Albert Ziegler, Eirini Kalliamvakou, Shawn Simister, Ganesh Sittampalam, Alice Li, Andrew Rice, Devon Rifkin, and Edward Aftandilian. 2022. Productivity assessment of neural code completion. (2022). arXiv: 2205.06537 [cs. SE].