

The Struggles of LLMs in Cross-Lingual Code Clone Detection

MICHELINE BÉNÉDICTE MOUMOULA*, CITADEL Excellence Centre in AI4D, Burkina Faso and University of Luxembourg, Luxembourg

ABDOUL KADER KABORÉ, University of Luxembourg, Luxembourg

JACQUES KLEIN, University of Luxembourg, Luxembourg

TEGAWENDÉ F. BISSYANDÉ, University of Luxembourg, Luxembourg

With the involvement of multiple programming languages in modern software development, cross-lingual code clone detection has gained traction within the software engineering community. Numerous studies have explored this topic, proposing various promising approaches. Inspired by the significant advances in machine learning in recent years, particularly Large Language Models (LLMs), which have demonstrated their ability to tackle various tasks, this paper revisits cross-lingual code clone detection. We evaluate the performance of five (05) LLMs and, eight prompts (08) for the identification of cross-lingual code clones. Additionally, we compare these results against two baseline methods. Finally, we evaluate a pre-trained embedding model to assess the effectiveness of the generated representations for classifying clone and non-clone pairs. The studies involving LLMs and Embedding models are evaluated using two widely used cross-lingual datasets, XLCoST and CodeNet.

Our results show that LLMs can achieve high F1 scores, up to 0.99, for straightforward programming examples. However, they not only perform less well on programs associated with complex programming challenges but also do not necessarily understand the meaning of “code clones” in a cross-lingual setting. We show that embedding models used to represent code fragments from different programming languages in the same representation space enable the training of a basic classifier that outperforms all LLMs by ~1 and ~20 percentage points on the XLCoST and CodeNet datasets, respectively. This finding suggests that, despite the apparent capabilities of LLMs, embeddings provided by embedding models offer suitable representations to achieve state-of-the-art performance in cross-lingual code clone detection.

CCS Concepts: • **Computing methodologies** → **Machine learning approaches**; • **Software and its engineering** → **Software maintenance tools**.

Additional Key Words and Phrases: Cross-Language Pairs, Code Clone Detection, Large Language Model, Prompt Engineering, Embedding Model

ACM Reference Format:

Micheline Bénédicte Moumoula, Abdoul Kader Kaboré, Jacques Klein, and Tegawendé F. Bissyandé. 2025. The Struggles of LLMs in Cross-Lingual Code Clone Detection. *Proc. ACM Softw. Eng.* 2, FSE, Article FSE047 (July 2025), 23 pages. <https://doi.org/10.1145/3715764>

*Corresponding author (micheline.moumoula@uni.lu)

Authors' Contact Information: [Micheline Bénédicte Moumoula](#), CITADEL Excellence Centre in AI4D, Ouagadougou, Burkina Faso and University of Luxembourg, Luxembourg, micheline.moumoula@uni.lu; [Abdoul Kader Kaboré](#), University of Luxembourg, Luxembourg, abdoulkader.kabore@uni.lu; [Jacques Klein](#), University of Luxembourg, Luxembourg, jacques.klein@uni.lu; [Tegawendé F. Bissyandé](#), University of Luxembourg, Luxembourg, tegawende.bissyande@uni.lu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 2994-970X/2025/7-ARTFSE047

<https://doi.org/10.1145/3715764>

1 Introduction

Copy-pasting a code fragment with or without change is common practice in software development. It leads to what is commonly referred to as code clones [34]. While clones provide some benefits (e.g., faster prototyping, reuse of proven code), they are also known to contribute to reducing the quality of the code, making maintenance costly, and being a source of bugs [17, 34]. Research shows that we can find between 5% to 23% of clones in a software system [22, 47]. The literature enumerates four types of clones [19]. Types 1, 2, and 3 clones are generally classified as syntactic clones, and type 4 is referred to as semantic clones. Syntactic clones occur when there are textual similarities, while semantic clones represent functional similarity. Type 4 is the most complicated clone and the hardest to detect [6].

Code clone retrieval is generally studied in the context of program fragments written with the same programming language. However, in modern software development, the multiplicity of software components involves various programming languages used, each for each efficiency in addressing specific programming challenges [21]. Developers may for example intentionally introduce clones for building identical systems for different platforms. Since the development of those systems is, for the most part, collaborative, if an expert in a certain type of programming language changes one part of the software for a given language, the exact functional change must be carried over to the other variants. This is more resource and time-consuming than modifying a software system based on a single language, as it requires prior knowledge of the system architecture and an understanding of the code changes performed by the first developer [28]. To manage cross-lingual systems in an easy, time-effective, and cost-effective way, developers need an automatic system that can detect clones of various languages at the same time [28].

The literature includes several approaches and tools for code clone detection. For instance, CCFinder [18], Deckard [16], NiCad [33], SourcererCC [38], NIL [29], and OREO [37] have been developed for type 1, 2, and 3 clones detection within the same programming language. For semantic clones within the same language, RtvNN [47], CDLH [45], DeepSim [50], and FA-AST [44] have been proposed.

Concerning cross-lingual clones detection, tools such as LICCA [43], CLCMiner [3], CLCDSA [28], OneSpace [9], C4 [40], TCCCD [10], CCT-Code [39], AdaCCD [7] and [20] have been introduced. They rely on deep learning techniques to capture the syntactic and semantic relationships between different parts of the source code. CLCDSA employed on-the-fly and deep neural network-based approaches using Abstract Syntax Trees (ASTs), a fundamental data structure in computer science for representing code structure, to extract features for cross-lingual code clone detection. On the fly approach measures cosine similarity between the two matrices generated by extracting values of the features from source code fragments while the deep neural network-based approach learns automatically the values of features from data and detects cross-lingual clones. OneSpace introduces a novel approach to cross-language code clone detection by mapping programming languages into a shared embedding space using both code and API data. It employs a Siamese network to evaluate the similarity between embedded programs and explores an alternative method using Word Mover's Distance for efficiency. The approach also examines the influence of factors such as code complexity and training data size on its performance. TCCCD employed the UniXcoder [12] pre-trained model to map code clones from different languages into vector space and fine-tuned the model using triplet learning. Sorokin et al. proposed an approach using cross-consistency training (CCT) for language models training on source code in different programming languages and presented the CCT-LM model, initialized with GraphCodeBERT and fine-tuned with CCT to retrieve semantically similar codes given a code as the query. AdaCCD leverages language-agnostic representations from

pre-training programming language models (GraphCodeBERT and CodeBERT) and contrastive learning for cross-lingual adaptation to detect code clones.

The rise in recent years of Large Language Models (LLMs) has led to the introduction of new studies in various domains, including software engineering. As regards code clone detection, the researcher also investigated the use of LLMs.

A representative work on this topic is proposed by Dou et al. [6], who proposed an approach with diverse prompt engineering techniques to assess LLMs' ability to detect clones in several programming languages. The experimental results showed that GPT-3.5-turbo and GPT-4¹ excel in detecting complex semantic clones within the same programming language, surpassing existing methods. However, this study only focuses on code clone detection in the same language.

In this paper, we investigate Large Language Models (LLMs), and Embedding Models' (EM) capabilities for cross-lingual code clone detection. LLMs performances being influenced by prompt engineering [2, 25, 49], we first design various different prompts that we applied on five LLMs (Falcon-7B-Instruct [1], LLAMA2-Chat-7B [41], Starchat- β [42], StarCoder2-15b-Instruct [46] and GPT-3.5-Turbo). The design prompts include two simple prompts: one that asks the LLMs to directly provide "yes/no" answers if the two code snippets are clones or not, and an improved version that includes instructions asking the LLMs to take into account the whole structure and logic of the two code snippets and conclude if they perform a similar task. In addition to those simple prompts, we designed one-step and multi-step Chain of Thought (CoT) [6] prompts to evaluate their impact on LLMs' performances. With these prompts, we want the LLMs to follow step-by-step reasoning to perform the task by taking into account the code similarity, the similar lines, the thinking process, and the code description analysis.

Our second study for cross-lingual code clone detection relies on the use of embedding models. We selected three pre-trained models: GraphCodeBERT [14] and UniXcoder [13], two widely used open-source embedding models for both natural language (e.g., descriptions, comments) and programming languages (source code), and the 'Text-embedding-3-large' provided by OpenAI.

For GraphCodeBERT and UniXcoder models, which are also used as baselines, similar to prior work [10, 20] we directly leverage their ability to be modified for downstream tasks (in our case, cross-lingual code clone detection).

For the 'Text-embedding-3-large' model, given two code snippets, we directly provide them as inputs to the embedding model, which outputs two vectors. To detect whether the provided code snippets are clones or not, we can directly perform a similarity measure using the two vectors and cosine similarity [31]. This requires the definition of a threshold. In our experiments, we investigate several threshold values. With the embedding vectors, we also conducted additional experiments for clone detection by training custom binary classifiers.

We perform evaluations using both the XLCOST [51] and CodeNet [32] datasets. We identified those datasets with respect to the large number of samples included and the various programming covered by covered languages. While the datasets have been used in the literature for many cross-lingual related tasks [11, 20, 23, 27], they are not directly usable as benchmarks for code clone detection: ① clone pairs must be established, and non-clone pairs must be carefully constructed. We undertake this effort and contribute to the community with a benchmark that includes a balanced set of 6 000 positive pairs (i.e., clones) and 6 000 negative pairs (i.e., non-clones).

Our experimental results show that GPT-3.5-Turbo achieved an overall F1 score of 0.99 for the XLCOST and 0.82 for CodeNet. It outperforms LLAMA2-Chat-7B, Starchat- β , StarCoder2-Instruct [46] and Falcon-7B-Instruct that achieve respectively overall F1 scores of 0.82, 0.615, 0.595, and 0.48 for the XLCOST dataset and 0.58, 0.52, 0.47, and 0.43 for the CodeNet dataset. Among all the

¹<https://openai.com/>

designed prompts, the “improved simple prompt” is the one that enables all LLMs to achieve their best performance in the task of code clone detection. This underlines that LLMs’ understanding of the notion of clones is narrow but that, once provided with definition guidance, they are relatively able to reason across programming languages.

Regarding the embedding model, we first note that the cosine similarity derived from code snippets’ vectors can be directly used to identify code clones. Indeed, with cross-lingual pairs, the embedding model produces close vectors when the code snippets are clones and distant vectors when they are non-clones. This is observed in $\sim 74\%$ of the pairs with a cosine similarity threshold of 50%. We further leverage the vector pairs to train several classifiers for cross-lingual code clone detection. Among those classifiers, the Support Vector Machine classifier (SVM) using a polynomial kernel achieves F1-scores of 1 and 0.986, respectively, for XLCoST and CodeNet. In sum, all our evaluations reveal that the embedding model outperforms all LLMs, even though the latter yields satisfactory results when combined with CoT-based prompts.

This research addresses the challenging problem of cross-language clone detection, which requires a deeper semantic understanding of code snippets across diverse syntaxes, structures, and programming styles. By leveraging the capabilities of large language models (LLMs), we explore their potential and limitations in this complex domain and propose strategies, including enhanced prompt designs, to improve their effectiveness. We introduce an approach that combines binary classifiers with code embeddings to effectively capture semantic similarities in cross-language scenarios. This method offers a practical solution for analyzing multilingual codebases, addressing critical challenges in areas such as multilingual code refactoring, security analysis, and software maintenance. By focusing on the application of LLMs to cross-language clone detection, this research contributes to an underexplored area with significant implications for managing large and diverse software ecosystems. This work makes several key contributions to the field of cross-lingual code clone detection:

- We conduct a large-scale investigation of LLMs’ capabilities for cross-lingual code clone detection. Our study considers Java paired against 10 programming languages, making it the largest study applying various LLMs to various cross-lingual datasets. Furthermore, compared to existing literature, we assess the influence of different prompt engineering techniques.
- We provide insightful findings on the performance of LLMs in code clone detection. Notably, we show that the similarity between two programming languages influences the ability of LLMs to detect code clones across them, mainly when the prompt is simple. If the prompt instructs the LLM to focus on logic/reasoning, programming language differences have substantially less influence. Eventually, we discuss the generalizability and overall effectiveness of LLMs for the task of cross-lingual code clone detection.
- We discuss the performance of LLMs by comparing them against two baselines and the classical approach of machine learning classification based on learned representations of code. Experimental results even hint that LLMs may not even “understand” the meaning of “clones” in the context of clones.

The remainder of this paper is organized as follows: [Section 2](#) presents background. [Section 3](#) introduces the experimental setup. [Section 4](#) evaluates our analysis. [Section 6](#) discusses related work. [Section 7](#) concludes.

2 Background

This section presents a short introduction to the concepts of code clone and LLMs.

2.1 Code Clones

A pair of code clones represents two code fragments that are similar, by some given definition of similarity [35]. We can enumerate four types of clones :

- Type-1 or identical code fragments represent the same code except for white space, comments, and layout.
- Type-2 or lexical code snippets represent identical clone pairs except for differences in variables or function names with Type-1 clone differences.
- Type-3 or syntactically represent similar code fragments that differ at the statement level. The code fragments differ in some lines with removed or added of some lines in addition to type-2 clone differences.
- Type-4 or semantic code clone represents code snippets that perform the same functionality but the implementation is different. In global they are syntactically dissimilar.
- Cross lingual clone is a semantic clone when code fragments are written in different programming languages. Our work is focused on semantic code clone detection for cross-lingual code snippets.

2.2 Large Language Models

LLMs are a type of artificial intelligence (AI) trained on massive datasets of text and code. This training enables them to perform a variety of natural language processing (NLP) tasks, including text summarization, text completion, machine translation, and even creative text generation. While LLMs are still under development, they have demonstrated the ability to generate human-quality responses to prompts.

Prompts are instructions or queries provided to an LLM to guide its output towards a specific task. Prompt engineering is a crucial aspect of LLM performance and refers to the optimization of prompts to achieve the most accurate and desirable outputs [36]. Chain-of-thought prompting is a specific technique that utilizes a sequence of connected instructions to guide the LLM's reasoning process step-by-step, resulting in demonstrably more accurate solutions for complex problems

3 Experimental Setup

First, we articulate the research questions that we aim to address. Second, we introduce the models that are leveraged in this study before presenting the datasets and describing the metrics. Finally, we introduce our experimental methodology on applying LLMs and embedding models for cross-lingual code clone detection.

3.1 Research Questions

We conduct our study around five research questions. The first four research questions are related to evaluating LLMs' performances in detecting cross-lingual code clones. In the last research question, we also evaluated the effectiveness of an embedding model and compared its performance against LLMs.

- **RQ1:** *To what extent does prompt engineering improve the effectiveness of LLMs in detecting cross-lingual code clones?* With this research question, we provide empirical results on the effectiveness of LLMs for the cross-lingual code clone detection task. The experiment comprehensively considers various prompting strategies to provide insights into the impact of prompt designs across several LLMs.
- **RQ2:** *To what extent do LLMs actually understand the task of code clone detection in a cross-lingual setting?* Code clone detection is about comparing the similarity of code fragments. In a cross-lingual setting, an LLM must make abstractions to achieve similarity comparisons. With this research question, we investigate how LLMs define "clones" and whether they can be guided

towards making the necessary abstractions for the task of code clone detection in a multi-lingual setting.

- **RQ3:** *How does the similarity between programming languages influence the performance of LLMs in cross-lingual code clone detection?* Because programming languages may or may not share similarities in terms of syntax, we hypothesize that LLMs performance may be affected by different configurations of the cross-lingual clone detection tasks.
- **RQ4:** *Are LLMs outperforming classification models for cross-lingual code clone detection?* Before the advent of LLMs, state-of-the-art approaches for code clone detection were employing machine learning-based classification, notably using deep representation learning. We thus compare LLMs against baseline classifiers (e.g. with SVM and kNN) based on Embedding models. The objective is to shed lights into the power of representations vs reasoning complexity in LLMs.

3.2 Models

We evaluated five LLMs and one embedding model for cross-lingual code clone detection. The selection process prioritized models fine-tuned for programming languages and code-related tasks, as well as those offering reproducibility and accessibility. From OpenAI, we included the popular GPT-3.5-Turbo model due to its prevalence in software engineering research and its balance of capability and cost-effectiveness, making it suitable for large-scale evaluations. To explore open-source options, we include Falcon-7B-Instruct, StarChat- β , and LLAMA2-Chat-7B. These models are trained on multiple programming languages and some have been employed in prior studies for code clone detection within the same programming language. Among these, StarCoder2-15b-Instruct was selected for its robust fine-tuning on programming languages, making it particularly well-suited for our task. Additionally, we incorporated the Text-embedding-3-large model, an embedding model from OpenAI's web API, to investigate the effectiveness of pre-trained embeddings for this task. By including this model, we extended our analysis beyond LLMs to assess the broader potential of embedding-based approaches. we provide detailed descriptions of each model in the following section.

- GPT-3.5-Turbo is an LLM developed by OpenAI, known for its strong performance in various language tasks. It excels at understanding and generating text in a conversational context, producing coherent and relevant responses based on user input. Notably, the model can retain information short-term, facilitating meaningful dialogue. Trained on a massive dataset of internet text, GPT-3.5-Turbo possesses a broad knowledge base and can leverage this information for diverse language processing applications.
- The Falcon-7B-Instruct model is a large language model (LLM) with 7 billion parameters. It is based on the Falcon-7B architecture and has been fine-tuned on a combined dataset of conversational and instructional text. This fine-tuning process optimizes the model for performing tasks typically associated with virtual assistants.
- The LLAMA2-Chat-7B model is a large language model (LLM) with 7 billion parameters. It is trained on a large corpus of text data, including conversational elements such as chat logs and social media posts. This training allows the model to acquire the patterns and structures of natural language dialogue, enabling it to generate coherent and contextually relevant responses to user prompts.
- StarChat- β is an instruction-tuned large language model (LLM) designed to act as a helpful coding assistant. The model is based on StarCoderPlus [24], a 15.5-billion parameter LLM trained on English and over 80 programming languages. StarChat- β is further fine-tuned on a specifically designed, "uncensored" variant of the open-assistant guanaco dataset. This fine-tuning process

strengthens the model's ability to understand and respond to human instructions within the context of coding tasks.

- StarCoder2-15b-Instruct [46] is the first code LLM trained entirely through a self-aligned process, eliminating the need for human annotations or distilled data from proprietary models. Our open-source pipeline leverages StarCoder2-15B [26] to generate a substantial dataset of instruction-response pairs. These pairs are then used to iteratively refine the model's capabilities, ensuring a fully transparent and permissive training process.
- Text-embedding-3-large is a text embedding model created by OpenAI. It is designed for tasks such as search, clustering, recommendations, anomaly detection, diversity measurement, classification text similarity evaluation, and code search. The model generates numerical representations (embeddings) of text or code inputs. Embeddings are low-dimensional, dense vector representations that capture the semantic relationships between words or code elements. This allows computers to more efficiently process and understand the meaning of text and code.

Baselines. We compare the results of the LLMs against two baseline models, chosen for their robust performance in previous research, to strengthen our evaluation.

- The GraphCodeBERT [15] is a transformer-based model that prioritizes data flow over traditional syntactic structures, such as abstract syntax trees (AST), in its pre-training phase. By leveraging structure-aware approaches, it enhances code representation, making it particularly effective for tasks like code clone detection, translation, refinement, and search. This focus on the intrinsic flow of data within code enables GraphCodeBERT to capture deeper semantic relationships, leading to improved performance across various code understanding tasks.
- The Unixcoder [12] is a cross-modal pre-trained model for programming languages that improves upon encoder-decoder frameworks by using mask attention matrices with prefix adapters for fine-grained control. It leverages data such as abstract syntax trees (ASTs) and code comments to enhance source code representation. Pre-trained on nine programming languages, including Java, Ruby, Python, PHP, Javascript, Go, C, C++, and C#, UniXcoder offers enhanced capabilities in representing cloned code, particularly in C++ and C#.

3.3 Datasets

We selected two cross-lingual benchmarks to ensure a robust and comprehensive evaluation across all considered programming languages: XLCoST [51] and CodeNet [32].

- The **XLCoST** dataset [51] is collected from GeeksForGeeks², a popular online platform for computer science education and skill development. Notably, it hosts coding challenges for programmers. The dataset includes samples from seven different programming languages (C, C++, C#, Java, JavaScript, PHP, Python). Since the dataset includes only true pairs of cross-lingual code clones, we must undertake to create negative clone pairs.
- The **CodeNet** dataset [32] comprises a large collection of code samples with extensive metadata across 4000 coding problems and over 50 programming languages (although C++, C, Python, and Java are the dominant languages). The code samples are annotated with a rich set of information, such as the problem ID and the status, which indicates acceptance or error types. The data is collected from AIZU Online Judge³ and AtCoder⁴. Based on the problem ID, we can construct both, positive and negative clone pairs

²<https://www.geeksforgeeks.org/>

³<https://onlinejudge.u-aizu.ac.jp/>

⁴<https://atcoder.jp/>

Construction of sample pairs. To ensure the highest quality samples for our experiments, we implemented a rigorous selection process for both non-clone and clone pairs. Both raw datasets include samples with a problem ID, description, and code snippets in various programming languages. This structure allows us to easily retrieve positive pairs, defined as “two code snippets in different programming languages addressing the same programming problem”. However, it is challenging to build a ground truth set of negative pairs. We postulate that a pair can be negative for two main reasons: (i) the problems addressed by the two code snippets are different; (ii) for the same problem two code snippets actually differ semantically (e.g., one is incorrect). For constructing our dataset, we only select code validated by developers. This means that negative pairs are only those where the problems are different. We therefore propose to use sentence embeddings to compare problem descriptions and build the dataset of negative pairs. Here is the breakdown:

- ❶ For each problem description, we generate the corresponding sentence vector using a state-of-the-art sentence vector generation model (all-mpnet-base-v2⁵). This model has been widely validated for its effectiveness in capturing semantic meaning.
- ❷ To group similar problems, we applied the DBSCAN clustering algorithm to the sentence vectors. Within each cluster, we selected the most complicated problem based on a pre-computed complexity metric. We identify the furthest problem (by vector distance) as a candidate for a negative pair (considering all the problems in the dataset).
- ❸ **Negative set selection:** We ensure that there are no duplicates or over-representation of specific problems and we prioritize pairing the most complex code.
- ❹ **Positive set selection:** The remaining problems were used to construct the positive set, with the most complex code selected to represent the clone pairs.
Before the previously described process, we calculated the cyclomatic complexity of each problem and code sample provided by the two datasets selecting the most complex instances to ensure a representative and challenging dataset.

For our experiments, we used subsets of 3000 clone pairs and 3000 non-clone pairs from each dataset, with Java selected as the primary language cloned alongside four and ten programming languages, respectively, for the XLCoST and, the CodeNet dataset with the same ratio for each programming language.

For the Baselines, We employ a similar selection process. However, we restricted our dataset to CodeNet, given its broader coverage of programming languages. From CodeNet, we selected only the programming languages that were compatible with the baseline models. Additionally, we filtered the dataset to include only pairs that met the input size requirements of the baseline models.

3.4 Metrics

Code clone detection can be considered as a binary classification problem. Therefore, we rely on classical metrics for classification evaluation, notably Precision, Recall, and F1 score. Precision represents the proportion of true, correct positive predictions to all of the positive predictions ($TP/(TP+FP)$). The recall represents the proportion of true correct positive prediction to all of the true positive samples ($TP/(TP+FN)$). F1 score is defined as the harmonic mean of precision and recall.

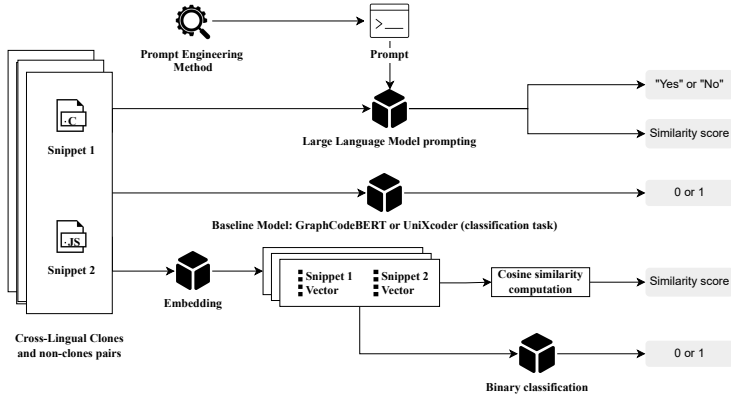


Fig. 1. Experimental workflow for cross-lingual code clone detection using LLMs vs Classification Models

3.5 Methodology

We overview the experimental methodology that is employed for assessing LLMs as well as baselines, and the classification models for the task of cross-lingual code clone detection. Figure 1 illustrates the process, highlighting the input datasets, the models as well as the expected outputs.

3.5.1 Cross-Lingual Code Clone Detection as an NLP Task.

We investigate the potential of LLMs for cross-lingual code clone detection using prompt engineering to explore the highest potential that LLMs can provide for this task.

We build on zero-shot and chain-of-thought prompting techniques and design eight prompts which, each include a specially-crafted instruction as well as the raw content of the pair of code fragments to analyse for clone detection. The expected output format varies depending on the prompt design: it can be either a straightforward binary "yes/no" answer or a numerical value representing a similarity score.

The proposed experimental methodology aims to assess the LLMs' ability for actual analysis of code semantics across different programming languages. Ultimately, we can provide measurements on the efficacy of LLMs for accurate cross-lingual code clone detection.

LLM Prompting. Our experiments are considering several prompts designed following two strategies: Zero-shot prompting, which is a classical and baseline approach used in the current literature [7] and Chain-of-Thought, which is a state-of-the-art prompting strategy that has been shown successful for several software engineering tasks [6]. All prompts are provided in Table 1.

3.5.2 Cross-Lingual Code Clone Detection as a Classification Task.

Text-embedding-3-large Model. We investigate, under the same conditions as LLMs, the potential for traditional classification mechanisms based on machine learning models. To that end, we follow the steps used by Keller et al. [19] in their learning-based approach for code clone detection. In their approach a classifier is trained based on sample sets of clone and non-clone pairs where each pair is represented by the absolute difference between the embedding vectors of each code fragment in the pair. Compared to other concatenation techniques that were assessed, their

⁵<https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

Table 1. List of Prompts Designed to Assess the Effectiveness of LLMs for the Task of Cross-Lingual Code Clone Detection

(a) Initial Pre-designed Prompts

Prompting	Name	Prompt	output
Zero-shot	Simple prompt	Analyze the following two code snippets and determine whether they are clones, regardless of the programming language. Respond with 'yes' if the code snippets are clones or 'no' if not.	yes/no
	Similar line	Analyze the following two code snippets for code clone detection, regardless of the programming language. You should first report which lines of code are more similar. Then based on the report, please answer whether these two codes are a clone pair. The response should be 'yes' or 'no'.	yes/no
	Reasoning	Provide a detailed reasoning process for detecting code clones in the following two code snippets, regardless of the programming language. Based on your analysis, respond with 'yes' if the code snippets are clones or 'no' if they are not.	yes/no
Chain of thought	Integrate	Analyze the following two code snippets to assess their similarity and determine if they are code clones, regardless of the programming language. Provide a similarity score between 0 and 10, where a higher score indicates more similarity. Additionally, presents a detailed reasoning process for detecting code clones. Conclude by 'yes' if they are clones or 'no' otherwise.	yes/no
	Separate code	Step 1 : Analyze the following code snippet and explain the function of the snippet. Step 2 : Analyze the following functions of two code snippets and determine if they are code clones, regardless of the programming language. The function of the first code snippet is {step 1 output} and the function of the second is {step 1 output}. Please answer 'yes' if the code snippets are clones, regardless of the programming language, or 'no' if they are not.	yes/no
	Separate Explanation	Step 1 : Similarity/Reasoning/Difference/Integrated process without the cloning conclusion. Step 2 : Analyze the following two code snippets and determine if they are code clones. The Clone Similarity/Reasoning/Difference Integrated information of the first and the second code is: Please respond with 'yes' if the code snippets are clones or 'no' if they are not.	yes/no
	Code similarity	Assess the similarity of the following two code snippets and provide a similarity score between 0 and 10. A higher score indicates that the two codes are more similar. Output the similarity score.	Similarity score

(b) Improved Prompt - Explicitly Describing the expected meaning of "Clone" to the LLM for cross-lingual clone detection

Prompting	Name	Prompt	output
Zero-shot	Improved simple prompt	Consider the overall structure and logic of the following two codes and determine if the two code snippets perform a similar task. Respond with 'yes' if the two codes perform similar tasks or 'no' otherwise.	yes/no

work conclude that differencing achieved the best results. In our case, however, the representations are yielded Embedding models that are provided alongside the LLM infrastructure of OpenAI. We use Text-embedding-3-large⁶ embedding model, which was trained on massive code and text data to yield representations that maximally capture semantics and can serve for similarity analysis. Finally, we train a binary classifier based on the yielded representations and using two basic learners: k Nearest Neighbors (k-NN) and Support Vector Machines (SVM).

Beyond the learning-based classification approach, we also investigate a straightforward approach where we compute the similarity between two fragments in a cross-lingual pair represented in a single representation space. To that end, we apply the cosine similarity metric on the Ada embeddings and check for the optimal threshold value to achieve the highest performance in terms of clone pair identification.

Baseline Models. We fine-tune the selected baseline models to optimize their performance for the task. These fine-tuned baselines are then systematically compared to the large language models (LLMs) to evaluate their relative effectiveness. This comparison allows us to assess how traditional models perform compared to the LLMs across the same task, providing insights into the strengths and limitations of each approach.

⁶<https://openai.com/index/new-and-improved-embedding-model/>

4 Evaluation

For each research question presented in 3.1, we first present the experiment we conducted and then show the result.

Table 2. Performance of GPT-3.5-Turbo with Various Prompts on the Task of Cross-lingual Code Clone Detection

Prompt name	Clone type	XLCoST			CodeNet		
		Recall	Precision	F1-score	Recall	Precision	F1-score
Simple prompt	Non-clone	0.99	0.97	0.98	1.00	0.56	0.72
	Clone	0.97	0.99	0.98	0.21	1.00	0.34
Khajezade et al. [20]	Non-clone	-	-	-	1.00	0.56	0.72
	Clone	-	-	-	0.21	1.00	0.35
Similar line	Non-clone	1.00	0.98	0.99	1.00	0.60	0.75
	Clone	0.98	1.00	0.99	0.33	0.99	0.49
Reasoning	Non-clone	0.98	0.99	0.98	0.99	0.71	0.82
	Clone	0.99	0.98	0.98	0.59	0.98	0.74
Integrate	Non-clone	0.99	0.90	0.94	0.99	0.60	0.74
	Clone	0.89	0.99	0.93	0.32	0.98	0.49
Separate code	Non-clone	1.00	0.98	0.99	1.00	0.58	0.73
	Clone	0.98	1.00	0.99	0.27	1.00	0.42
Separate explanations	Non-clone	0.98	0.99	0.98	0.99	0.73	0.84
	Clone	0.99	0.98	0.98	0.64	0.99	0.78

Table 3. Performance Comparison of LLMs on the task of Cross-Lingual Code Clone Detection - based on the "Simple Prompt"

LLM	Clone type	XLCoST			CodeNet		
		Recall	Precision	F1-score	Recall	Precision	F1-score
Falcon-Instruct-7B	Non-clone	0.15	0.51	0.23	0.11	0.48	0.17
	Clone	0.86	0.50	0.63	0.89	0.50	0.64
StarChat-Beta-16B	Non-clone	0.62	0.47	0.53	0.76	0.49	0.60
	Clone	0.30	0.45	0.36	0.22	0.48	0.30
LLaMA2-Chat-7B	Non-clone	0.01	1.00	0.01	0.04	0.68	0.08
	Clone	1.00	0.50	0.67	0.98	0.51	0.67
StarCoder2-15b-instruct	Non-clone	1.00	0.5	0.67	1.00	0.5	0.67
	Clone	0.01	1.00	0.03	0.02	0.96	0.03
GPT-3.5-Turbo	Non-clone	0.99	0.97	0.98	1.00	0.56	0.72
	Clone	0.97	0.99	0.98	0.21	1.00	0.34

4.1 [RQ.1] LLM Performance on Cross-Lingual Code Clone Detection and Impact of Prompt Engineering


Goal. We aim to fill a gap in the literature with a comprehensive analysis of the effectiveness of LLMs on the task of cross-lingual code clone detection. In particular, we investigate the impact of state-of-the-art prompt engineering strategy on the performance of the LLMs: we thus experiment with Zero-shot and Chain-of-Thought (CoT)-based prompts. Depending on the desired output formats, the designed prompts fall into two categories:

- output is a *score* [0-10]: We design a prompt that instructs the models to generate a similarity score between 0 and 10 when given a pair of code fragments: the lowest score (0) means that the model finds no similarity between the fragments, while the highest score (10) means that the model finds that the similarity is perfect. This prompt is designed using the CoT prompting strategy.

- output is Yes/No: We design several prompts to guide the model towards answering with "Yes" or "No" when given a pair of code fragments. One of these prompts is based on a straightforward Zero-Shot strategy, while we provide five (5) prompts building on Chain-of-Thought. The design of the prompts is varied to take into account the requested depth of analysis (line by line or overall block), the specification of need to provide a reasoning of the process, the breaking of the steps for decision, etc. This methodology is inspired by prior work [6] leveraging LLMs for clone detection, which has not assessed cross-lingual capabilities.

Experiments. We evaluated the five LLMs considered in this study (cf. Section 3.2), namely GPT-3.5-Turbo, Starchat- β -16B, Falcon-7B-Instruct, LLAMA2-Chat-7B, and Starcoder2-15b-instruct on the constructed subsets (Section 3.3) from XLCoST and CodeNet. We use the ChatGPT API to prompt the GPT-3.5-Turbo LLM and parse the outputs to compute the precision, recall, and F1 scores finally. For prompts leading to outputs with a similarity score, we analyzed the LLMs' outputs across a range of threshold values for code clone decisions. Our experiments included a comparison of our prompts with a prior proposal for cross-language code clone detection, evaluated on our dataset.

Results. Our experiments confirmed that many LLMs actually fail to support the Chain-of-Thought prompting strategy, confirming a claim by Dou et al. [6]. Consequently, to analyze the performance of LLMs across different prompts, we focus on GPT-3.5-Turbo, which produced results when instructed with all our designed prompts. Table 2 summarizes the results for this LLM across all the "yes/no" output prompts. On the XLCoST dataset, the CoT-based "Separate explanations" prompt enables the model to achieve the highest F1 score (98% for non-clones and clones). On the CodeNet dataset, the same prompt achieves 84% and 78% respectively for non-clones and clones. We can observe that the performance on the CodeNet dataset is lower than that of the XLCoST dataset. This discrepancy in performance is due to the difference in complexity between the two datasets. Indeed, XLCoST is mainly composed of source code implementing simple problems such as "Compute modulus division by a power" while CodeNet is built based on coding challenges of higher complexity such as "Given a string, find the minimum number of swaps needed to rearrange the characters into a palindrome". Furthermore, some LLMs exhibit a clear difference in their ability to detect code clones compared to non-clones. For instance, on the CodeNet dataset, the GPT-3.5-turbo model shows a 35 percentage point variation in the f1 score (72% on non-clone vs. 37% on clones). On the same dataset, we also observe a high variation of 47 for the 'Falcon-Instruct-7B' (17% on non-clones vs. 64% on clones) model. Manual analysis revealed a tendency for this model to classify nearly all input code snippet pairs as clones, resulting in an inflated number of false positives.

[RQ-1.1]  In general, for simple and recurrent programming tasks, LLMs show a high ability to detect cross-lingual clones (e.g., XLCoST samples). However, when applied to code samples of challenging programming tasks (e.g., CodeNet), LLMs performance drop by 20% on average in terms of F1 score. This result is insightful as it calls for a different perspective in benchmarking cross-lingual code clone detection approaches. Indeed, while XLCoST has been built for such purpose, CodeNet appears to offer a more realistic setting, in terms of complexity for assessment.

The "Simple prompt" exhibits a relatively consistent performance across both datasets, suggesting its potential for broader applicability. A detailed performance comparison of all five LLMs using this prompt is presented in Table 3. Notably, the GPT-3.5-turbo model outperformed other LLMs in this analysis. Also, our experiments demonstrate that our proposed simple prompt achieved comparable results to the previously suggested one.

[RQ-1.2] Compared to Zero-Short prompting, which is mainly used in the literature, Chain of Thought prompting improves the performance of LLMs by 1-28 percentage points in terms of overall F1-score. It is further noteworthy that GPT-3.5-Turbo substantially outperforms the other studied LLMs in the task of cross-lingual code clone detection.

In Figure 2, we present the result of the GPT-3.5-turbo LLM with the "Code Similarity" prompt, designed to generate a similarity score between code snippets (0-10 scale). The results show that with a threshold of 5, the LLM achieves an overall F1 score of 0.98 and 0.78 on the XLCoST and the CodeNet datasets, respectively.

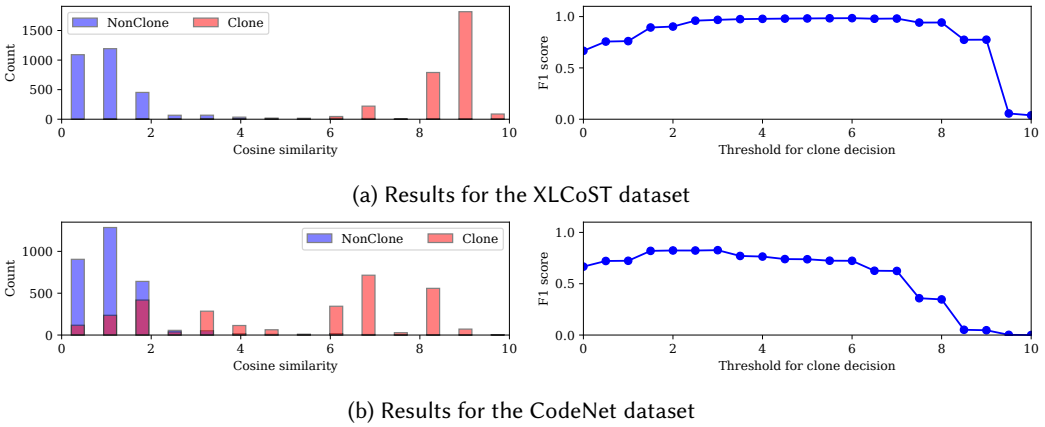


Fig. 2. Performance of GPT-3.5-Turbo prediction of Similarity Scores on the task of Cross-lingual Code Clone Detection

[RQ-1.3] Our experiments on the output requirements (Yes/No vs Similarity scores) underline that both approaches lead to similar performance scores. This suggests that the LLMs indeed consider the clone detection problem as a similarity computation task.

4.2 [RQ2] LLMs' Reasoning for Cross-Lingual Code Clone Detection

Goal. Our investigation of RQ.1 on the performance of LLMs for the cross-lingual code clone detection task has revealed, based on a manual analysis of some sample cases, that the code pairs are often obviously (for a human) a clone, and yet the LLM fails to identify it. In other cases, the pair is obviously (for a human) not a clone, and yet the LLM concludes that it is. We thus now aim to further study the extent to which the LLMs understand the task. Based on our analysis, notably of whether the LLMs understand what "clone" means in a cross-lingual setting, we propose an improved version of the Zero-shot prompt to validate that LLMs performance can be improved with a hint on the definition of clone.

Experiments. We qualitatively analyzed the outputs of all the experiments conducted in the first research question to collect insights in the failures of the LLMs for the task of cross-lingual code clone detection. Taking into account those results, we improve the Zero-shot prompt design and repeat the same experiments as in Section 4.1 based. In this RQ, we focus on the Zero-shot prompting, since it is a baseline, in order to better evaluate the impact of the improvement in the definition.

Results. Our qualitative review of LLMs' outputs reveals that Falcon-Instruct-7B and LLAMA2-Chat-7B consider, most of the time, any input pair of code fragments as clones, leading to an excessive amount of false positives. Conversely, Starcoder2-15b-instruct and Starchat- β considers most of the pairs as non-clones, yielding a significant number of false negatives.

The behavior of Starcoder2-15b-instruct and Starchat- β are explained by the fact the reasoning of the LLM completely breaks down for the task of cross-lingual code clone detection. The LLM even goes far as to over-confidently state that *two code snippets of different programming languages cannot be clones*. Similar breakdowns in LLM reasoning have recently been documented in the literature [30]. Our experiments reveal for future research development in LLMs that cross-lingual code clone detection is a relevant and easy-to-assess task for checking potential reasoning breakdown of LLMs. Unfortunately, Starcoder2-15b-instruct and Starchat- β are not the only LLM with a breakdown in reasoning: Similar situations are also apparent in several outputs of GPT-3.5-turbo.

We analyze specifically the cases where the prompts require the outputs to include explanations from the Chain of Thought. A major finding was that in most of the cases where the LLM output was accurate, the LLMs actually stated that their decision was made taking into the “**overall structure and logic**” of the code fragments. Based on this observation, we propose to design a new prompt, the ‘improved simple prompt,’ that instructs the LLMs on what to consider as code clones. In other words, the improved prompt clearly defines the concept of code clones without using the term, which may be misleading to the LLMs: we instruct the model, in a Zero-Shot manner, to consider the overall structure and logic of the code snippets independently of their programming language (cf. Table 1b).

Table 4 presents the performance achieved by the various LLMs based on the improved prompt. The performance improvement (compared to results when using a prompt that simply asks about “clones” - cf. Table 3) is substantial. In particular, for Starchat- β and LLAMA2-Chat-7B, the new prompt helps the LLM to exhibit an improvement between 7 and 48 percentage points in terms of F1 score.

[RQ-2] 🐼 *LLMs consistently are not fully aware of the notion of code clones in a cross-lingual setting. When prompted to adhere a definition of cloning as a measurement of similarity taking into account overall structure and logic, LLMs are able to demonstrate substantial success in the task of cross-lingual code clone detection. On the XLCOST dataset, where cross-lingual clones are known (cf. RQ1) to be easier to detect, the improved prompt enables the LLaMA2-Chat-7B model to raise its F1 performance score from 0.01 to 0.80 for non-clones and 0.67 to 0.83 for clones, reducing the gap from GPT-3.5-Turbo by 14.5 percentage points.*

4.3 [RQ3] Influence of the Programming Languages Syntactical Similarity on LLMs Performances

Goal. In the studied task, a clone pair involves two fragments from two distinct programming languages. Previous research questions focused on highlight the overall performance of LLMs on built pairs. We now attempt to check whether the specific combinations of programming languages may affect the performance of the LLM in identifying clones. In particular, we investigate whether the syntactic similarity of programming languages has an influence on the effectiveness of the LLMs.

Experiments. Some programming languages share common core concepts, such as for controlling program flow (e.g., loops, conditionals, etc.) or storing data (e.g., variables). In terms of Syntax, we note for example that Java's syntax is most like C++ and C. In fact, Java borrows heavily from C++ syntax, and therefore their programs may present similar structures. Similarly, languages like

Table 4. LLMs performance comparison with the improved prompt, designed based on LLMs common behavior

		XLCoST			CodeNet		
LLM	Clone type	Recall	Precision	F1-score	Recall	Precision	F1-score
Falcon-Instruct-7B	Non-clone Clone	0.23	0.54	0.32	0.15	0.51	0.23
		0.80	0.51	0.64	0.86	0.50	0.63
StarChat-Beta-16B	Non-clone Clone	0.66	0.60	0.63	0.75	0.52	0.61
		0.57	0.64	0.60	0.34	0.59	0.43
Starcoder2-15b-instruct	Non-clone Clone	0.96	0.58	0.73	0.92	0.52	0.67
		0.31	0.89	0.46	0.17	0.67	0.27
LLaMA2-Chat-7B	Non-clone Clone	0.75	0.87	0.81	0.78	0.57	0.66
		0.89	0.78	0.83	0.40	0.65	0.50
GPT-3.5-Turbo	Non-clone Clone	0.98	0.99	0.98	1.00	0.75	0.85
		0.99	0.98	0.98	0.66	1.00	0.79

Table 5. Performance of GPT-3.5-turbo on the task of cross-lingual code clone detection - Detailed F1 scores by programming language pairs (Java - Lang-X)


Lang-X	Simple Prompt	Improved Simple Prompt	Similar Line	Reasoning	Integrate	Separate Explanation	Separate Code
C	0.68	0.85	0.70	0.83	0.71	0.84	0.60
C#	0.54	0.86	0.65	0.81	0.66	0.82	0.60
C++	0.59	0.81	0.66	0.81	0.67	0.85	0.56
Go	0.46	0.81	0.57	0.74	0.61	0.82	0.54
Ocaml	0.48	0.81	0.59	0.73	0.54	0.75	0.56
PHP	0.50	0.78	0.62	0.80	0.64	0.82	0.59
Python	0.52	0.80	0.59	0.76	0.61	0.80	0.59
Ruby	0.49	0.81	0.59	0.78	0.52	0.76	0.57
Rust	0.47	0.86	0.57	0.75	0.57	0.80	0.55
JavaScript	0.53	0.85	0.63	0.80	0.63	0.82	0.56

C# (pronounced C-Sharp) share similarities with Java because they are object-oriented languages with similar features like classes, objects, inheritance, and encapsulation. To better highlight the potential combinations, we focus on the case of Java, and consider sets of pairs where the second programming language is varied across C, C#, C++, Go, Ocaml, PHP, Python, Ruby, Rust, and JavaScript. We then disaggregate the performance metrics by programming language pair, focusing on experimental results with the CodeNet dataset from RQ-1 (cf. [Section 4.1](#)). Given that we investigate the performance across the various prompt designs, we focus on the GPT-3.5-Turbo LLM, for which all prompt designs successfully returned answers.


Results. [Table 5](#) shows all the results across various prompts and by programming language pair. We note that the performance difference is at least of ~10 percentage points in terms of F1 score when the LLM is applied to fragments of Java-C language pair vs fragments of Java-Ocaml language pair. C is indeed very close to Java in terms of syntax and core concepts, while Ocaml employs different programming paradigms.

Nevertheless, the results also show that, while syntactic similarity of programming languages influences the performance of LLMs when instructed with a simple prompt, complex prompts with reasoning instructions (e.g., Reasoning) and the Improved simple prompt (which guides the LLM to consider logic) manages to reduce the performance gap for language pairs that are supposedly relatively different (e.g., Java-PHP).

[RQ-3] *For the task of cross-lingual code clone detection, LLMs, when instructed with simple zero-shot prompts that do not clarify the definition of clones, see their performance influenced by the similarity of the languages of the code fragments in a pair. However, when the prompt explores the reasoning potential of the LLM, the differences in the programming language become less critical.*




(a) The performance of Text-embedding-3-large mixed with cosine similarity on XLCosT dataset




(b) The performance of Text-embedding-3-large mixed with cosine similarity on CodeNet dataset

Fig. 3. In (a) and (b), we first measure the performance of Text-embedding-3-large mixed with cosine similarity at different thresholds for each dataset. Using a threshold of 0.5, we achieved an F1-score of 0.98 on the XLCosT dataset and 0.74 on the Codenet dataset.



(a) Performance of the learning-based classifiers on the XLCosT dataset.



(b) Performance of the learning-based classifiers on the CodeNet dataset

Fig. 4. Impact of the binary classification on cross-lingual code clone detection using code embedding

4.4 [RQ4] Binary Classifier vs. Baselines vs. LLMs

Goal. The primary objective of this research is to evaluate the performance of Embedding models in detecting cross-lingual code clones and to compare their effectiveness against large language models (LLMs). Code clone detection has been studied in the literature as a binary classification problem. Various techniques leveraging machine learning has then been adopted and showed promising results. We compare LLMs against implementations of such techniques. Since we aim to assess the added value of LLMs (beyond the representations), we build on the embedding model that is provided alongside the LLM by OpenAI.


Experiments. We rely on the identified embedding model (Text-embedding-3-large - cf. Section 3.2) to generate the representations of code fragments in our constructed benchmark. We conduct three separate experiments:

- (1) *straightforward similarity computation*: we use cosine similarity [31] to compute the distance between the embeddings of two code fragments in a given pair. We can identify an optimal threshold for code clone identification by considering varying threshold values. In this study, the threshold value is set to 5 for all evaluations.
- (2) *Baselines model training for downstream classification task*: We conducted a series of experiments using two baseline models (GraphCodeBERT and UniXcoder). Given the input limitations of the baseline models, we made some adjustments to overcome the input length constraints of the baseline models when processing the dataset. To ensure robust evaluation, the selected dataset was split into an 80% training set and a 20% validation set. This experimental setup allowed us to fine-tune the baseline models and assess their ability to identify cross-lingual code clones under comparable conditions to the LLMs.
- (3) *binary classifier training*: we consider some basic learners that we apply to the concatenation⁷ of embeddings in code fragment pairs to train a binary classification model. As learners, we consider Support Vector Machines (SVM) [4] and k-Nearest Neighbors (k-NN) [5]. For SVM, we use three variants with `ball_tree`, `kd_tree`, and `brute` algorithms, while for k-NN, we use four variants using the `linear`, `polynomial`, and `rbf` kernels. The experiments are performed using 10-Fold cross-validation on each dataset.

Results. Figure 3 describes the performance (F1 score) that can be achieved when using cosine similarity thresholds for deciding on code clones. The threshold set at 5 permits to achieve up to 0.99, 0.97, and 0.98 scores regarding the overall recall, precision, and F1 for samples associated with the XLCoST dataset. For the CodeNet dataset, the scores are 0.79, 0.84, and 0.78 for the overall recall, precision, and F1, respectively.

Table 6 show the performance of the baseline models. They demonstrated superior performance compared to the LLMs, achieving an improvement of 15 percentage points in terms of the overall F1-score.

Figure 4 summarizes the evaluation results of the trained classifiers. SVM achieves the highest scores at 1.00 for precision, recall, and F1. We further observe that k-NN achieves similar performance to SVM on the CodeNet dataset. This suggests homogeneity in the code samples enabling the learner to have a clear decision boundary. In contrast, k-NN lags well behind SVM on the XLCoST dataset, suggesting higher diversity of samples (albeit potentially simple programs).


[RQ-4.1]  The Text-embedding-3-large embedding model yields comprehensive representations of cross-lingual code fragments that enable rapidly identifying code clones based on straightforward similarity computation or learned classification.

⁷cf. Section 3.5 - following prior work [19], we use the Absolute difference of embeddings as a concatenation method

Table 6. Comparison of LLM, Baselines, and Binary Classifier Performance

Models	CodeNet				
	Name	Clone type	Recall	Precision	F1-score
Baselines	GraphCodeBERT	Non-clone	0.96	0.97	0.97
		Clone	0.98	0.97	0.97
		Average	0.97	0.97	0.97
	UnixCoder	Non-clone	0.94	0.88	0.91
		Clone	0.89	0.95	0.92
		Average	0.92	0.92	0.92
LLM	GPT-3.5-Turbo	Non-clone	1.00	0.75	0.85
		Clone	0.66	1.00	0.79
		Average	0.83	0.88	0.82
Binary classifier	Text-embedding-3-large + SVM(kernel='poly')	Non-clone	0.99	0.98	0.99
		Clone	0.99	0.99	0.99
		Average	0.99	0.98	0.99

To provide a comprehensive understanding of our results, we present in [table 6](#) an in-depth comparison between traditional baselines, LLMs, and binary classifiers.

[RQ-4.2]  Compared to the performance achieved by baselines (GraphCodeBERT and UnixCoder), and LLMs instructed by prompts of varying complexity, we find that simple models (binary classifiers) achieve better results. This suggests that a major challenge in cross-lingual code clone detection is identifying a single representation space for all programming languages rather than attempting to have a universal reasoning engine.

5 Discussion

This section explores the implications of our findings for cross-language code clone detection. Our investigation into the ability of LLMs to interpret cross-language code clones yielded valuable insights.

LLM Interpretation and Performance. Our experiments showed that LLMs can perform substantially better when we give them specific instructions about what to look for in the code (e.g., with the improved simple prompt) compared to when the prompt is generic and zero-shot based. The improved prompt, which instructs the LLM to reason about the overall structure and logic of the code leads LLMs to identify a large number of cross-lingual code clones (cf. [Table 4](#)). This finding suggests that LLMs can avoid dramatic reasoning failures when the prompt is sufficiently informative on the task: LLMs may need some guidance to see the bigger picture.

Representations vs. Prompt Engineering. The superiority of the learned classifier based on embedding models highlights the value of representations. Unlike prompts, which require careful design, an embedding model is a powerful, straightforward tool. It produces representations that capture the code's semantic meaning and functionality across different programming languages. Then based on the similarity of representations in a training set, a simple algorithm can learn decision boundaries, effectively overcoming language barriers in code clone detection. In contrast, prompts adds to the complexity and opacity of the reasoning workflow performed in an LLM.

Potential Data Leakage. Our experimental results using XLCOST demonstrate that the GPT-3.5-turbo exhibits a notable degree of familiarity with the samples. This may suggest potential leakage within this LLM training set.

Limitations and Threats to Validity. ① This study evaluated only five LLMs and three embedding models. Expanding the research to include a broader range of models with different architectures and training data could provide a more comprehensive understanding of LLM effectiveness in this task. ② The data used in our evaluations are derived from widely used datasets in the literature.

However, they may lack diversity, and this can lead to different observations on more diverse data (in terms of programming languages, complexity, ...). ③ Further exploration of prompt design strategies is needed. This could involve investigating variations of Prompt 2, exploring entirely different prompt structures, and leveraging insights from this study to refine prompts for more optimal LLM performance. ④ Due to limitations in current LLM capabilities, most prompt approaches did not always provide definitive "yes/no" responses for clone detection. This necessitated manual verification and correction of results, which constitute an internal threat to validity. Future research can focus on developing prompt design or LLM training strategies that yield more reliable and automated outputs. ⑤ The Text-embedding-3-large model is likely based on an earlier architecture of the GPT family. Newer text embedding models may help achieve even better performance. ⑥ The dataset used for fine-tuning GraphCodeBERT and UniXcoder was selected to align with both the input length constraints and the supported programming languages of these models. Consequently, the results reported in this work are limited to cases where input lengths do not exceed the maximum input size supported by the models, and the programming languages fall within the specific set supported by each model. ⑦ This study assumes a balanced dataset with an equal number of positive and negative pairs, which does not align with real-world conditions. In practical software projects, the vast majority of method pairs are not clones, resulting in a highly imbalanced scenario. While the balanced dataset used in this study provides insights into the techniques' effectiveness under controlled conditions, it may not fully represent their performance in real-world contexts. Future work should evaluate the studied techniques on datasets with more realistic imbalances to better understand their behavior and effectiveness in such scenarios. ⑧ The construction of negative pairs in the evaluation dataset introduces a potential threat to validity. While the authors employ an approach by distinguishing at the problem level to identify negative pairs, this design choice may inadvertently exclude certain types of negative pairs, particularly "close misses" pairs of code snippets that are similar but originate from distinct problems. As a result, the dataset may not fully capture the complexity of real-world cross-lingual clone detection scenarios, where such challenging distinctions are common. Future work could address this by exploring alternative sampling strategies to include a broader variety of "close misses" negative pairs, and by examining how different definitions of "negative pairs" influence the evaluation outcomes. This would help create a dataset that more comprehensively reflects the diversity of challenges in real-world settings.

6 Related Work

Prompt engineering is a challenging yet essential endeavor aimed at enhancing the performance of LLMs for specific tasks [48]. Ekin et al. [8] have demonstrated recently that to obtain accurate, relevant, and coherent output, users should design, refine, and optimize the prompt. Model understanding, domain specificity, and clarity are essential to identify the strengths and limits of the LLM, avoid ambiguity, and help it generate more accurate and relevant outputs. Our experimental results with a quick improvement of simple zero-shot prompts by providing a domain definition of the concept of clone confirms this finding for the field of software engineering.

Regarding code clone detection with LLMs, we can principally enumerate works by Dou et al. [6] and Khajezade et al. [20]. In [6], the authors proposed an automated code clone detection for all types of clones with various prompt designs for a comprehensive empirical evaluation of several LLMs applied to clones from a single language (Java). In alignment with our findings, they show that LLMs of the GPT family outperform others and that introducing intermediate reasoning steps through a chain of thought improves the performance of LLMs for the clone detection task.

Khajezade et al. [20] worked on mono and cross-lingual code clone detection using only GPT-3.5-Turbo (in contrast to our work, which explored various LLMs). They also used only zero-shot

prompts and considered only Java-Ruby combinations (in contrast to our work, which considered 10 language pairs). Their study furthermore does not investigate the power of representations over prompts.

7 Conclusion

This study explored the potential of LLMs for cross-language code clone detection. We evaluated five LLMs and investigated how prompt engineering enhances their performance. Additionally, we examined the effectiveness of code representations in the identification of code clones. Overall, our findings suggest the potential of using powerful LLMs for cross-language code clone detection. Nevertheless, While LLMs show promise, our results suggest that simpler binary classifiers leveraging robust language-agnostic representations can achieve higher, or at least comparable, performance for cross-language code clone detection. The insights gained from this study provide a foundation for future research aimed at developing more robust LLM-based techniques and improving code representation methods to advance software engineering practices. Furthermore, the methodologies and approaches presented here offer valuable reference points for continued exploration in this emerging domain. Future work could focus on addressing dataset imbalances by simulating "close misses" or exploring alternative sampling strategies to construct more representative negative pairs, thereby enhancing the reliability and applicability of cross-language code clone detection techniques.

Data Availability

To promote transparency and facilitate reproducibility, we make our artifacts available to the community at:

<https://github.com/TruX-DTF/CLCCD>

The repository includes the constructed benchmark of cross-lingual code clones, the prompt details, the experiment scripts, and the results.

Acknowledgments

This work was supported by (1) the Luxembourg Ministry of Foreign and European Affairs through their Digital4Development (D4D) portfolio under project LuxWAYS and (2) the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Project NATURAL - grant agreement N° 949014).

References

- [1] Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Merouane Debbah, Etienne Goffinet, Daniel Heslow, Julien Launay, Quentin Malartic, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. Falcon-40B: an open large language model with state-of-the-art performance. (2023).
- [2] Weiheng Bai, Qiushi Wu, Kefu Wu, and Kangjie Lu. [n. d.]. Exploring the Influence of Prompts in LLMs for Security-Related Tasks. ([n. d.]).
- [3] Xiao Cheng, Zhiming Peng, Lingxiao Jiang, Hao Zhong, Haibo Yu, and Jianjun Zhao. 2017. *CLCMiner: Detecting Cross-Language Clones without Intermediates*. *IEICE Trans. Inf. & Syst.* E100.D, 2 (2017), 273–284. doi:10.1587/transinf.2016EDP7334
- [4] Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine learning* 20 (1995), 273–297.
- [5] Thomas Cover and Peter Hart. 1967. Nearest neighbor pattern classification. *IEEE transactions on information theory* 13, 1 (1967), 21–27.
- [6] Shihan Dou, Junjie Shan, Haoxiang Jia, Wenhao Deng, Zhiheng Xi, Wei He, Yueming Wu, Tao Gui, Yang Liu, and Xuanjing Huang. 2023. Towards Understanding the Capability of Large Language Models on Code Clone Detection: A Survey. <http://arxiv.org/abs/2308.01191>
- [7] Yangkai Du, Tengfei Wu, Lingfei Wu, Xuhong Zhang, and Shouling Ji. 2024. AdaCCD: Adaptive Semantic Contrasts Discovery Based Cross Lingual Adaptation for Code Clone Detection. <http://arxiv.org/abs/2311.07277>

- [8] Sabit Ekin. 2023. Prompt Engineering For ChatGPT: A Quick Guide To Techniques, Tips, And Best Practices. doi:10.36227/techrxiv.22683919.v2
- [9] Mohammed El Arnaoty and Francisco Servant. 2024. OneSpace: Detecting cross-language clones by learning a common embedding space. *Journal of Systems and Software* 208 (2024), 111911.
- [10] Yong Fang, Fangzheng Zhou, Yijia Xu, and Zhonglin Liu. 2023. TCCCD: Triplet-Based Cross-Language Code Clone Detection. *Applied Sciences* 13, 21 (Nov. 2023), 12084. doi:10.3390/app132112084
- [11] Alex Gu, Baptiste Rozière, Hugh Leather, Armando Solar-Lezama, Gabriel Synnaeve, and Sida I Wang. 2024. Cruxeval: A benchmark for code reasoning, understanding and execution. *arXiv preprint arXiv:2401.03065* (2024).
- [12] Daya Guo, Shuai Lu, Nan Duan, Yanlin Wang, Ming Zhou, and Jian Yin. 2022. UniXcoder: Unified Cross-Modal Pre-training for Code Representation. <http://arxiv.org/abs/2203.03850>
- [13] Daya Guo, Shuai Lu, Nan Duan, Yanlin Wang, Ming Zhou, and Jian Yin. 2022. Unixcoder: Unified cross-modal pre-training for code representation. *arXiv preprint arXiv:2203.03850* (2022).
- [14] Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, Shujie Liu, Long Zhou, Nan Duan, Alexey Svyatkovskiy, Shengyu Fu, et al. 2020. Graphcodebert: Pre-training code representations with data flow. *arXiv preprint arXiv:2009.08366* (2020).
- [15] Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, Shujie Liu, Long Zhou, Nan Duan, Alexey Svyatkovskiy, Shengyu Fu, Michele Tufano, Shao Kun Deng, Colin Clement, Dawn Drain, Neel Sundaresan, Jian Yin, Daxin Jiang, and Ming Zhou. 2021. GraphCodeBERT: Pre-training Code Representations with Data Flow. <http://arxiv.org/abs/2009.08366>
- [16] Lingxiao Jiang, Ghassan Misherghi, Zhendong Su, and Stephane Glondou. 2007. DECKARD: Scalable and Accurate Tree-Based Detection of Code Clones. In *29th International Conference on Software Engineering (ICSE'07)*. IEEE, Minneapolis, MN, USA, 96–105. doi:10.1109/ICSE.2007.30
- [17] Elmar Juergens, Florian Deissenboeck, Benjamin Hummel, and Stefan Wagner. 2009. Do code clones matter?. In *2009 IEEE 31st International Conference on Software Engineering*. IEEE, Vancouver, BC, Canada, 485–495. doi:10.1109/ICSE.2009.5070547
- [18] T. Kamiya, S. Kusumoto, and K. Inoue. 2002. CCFinder: a multilingual token-based code clone detection system for large scale source code. *IEEE Trans. Software Eng.* 28, 7 (July 2002), 654–670. doi:10.1109/TSE.2002.1019480
- [19] Patrick Keller, Abdoul Kader Kaboré, Laura Plein, Jacques Klein, Yves Le Traon, and Tegawendé F. Bissyandé. 2022. What You See is What it Means! Semantic Representation Learning of Code based on Visualization and Transfer Learning. *ACM Trans. Softw. Eng. Methodol.* 31, 2 (April 2022), 1–34. doi:10.1145/3485135
- [20] Mohamad Khajezade, Jie JW Wu, Fatemeh Hendijani Fard, Gema Rodríguez-Pérez, and Mohamed Sami Shehata. 2024. Investigating the Efficacy of Large Language Models for Code Clone Detection. <http://arxiv.org/abs/2401.13802>
- [21] Pavneet Singh Kochhar, Dinusha Wijedasa, and David Lo. 2016. A Large Scale Study of Multiple Programming Languages and Code Quality. In *2016 IEEE 23rd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*. IEEE, Suita, 563–573. doi:10.1109/SANER.2016.112
- [22] Rainer Koschke. [n. d.]. Survey of Research on Software Clones. ([n. d.]).
- [23] Jia Li, Chongyang Tao, Zhi Jin, Fang Liu, and Ge Li. 2023. ZC 3: Zero-Shot Cross-Language Code Clone Detection. In *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 875–887.
- [24] Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, Qian Liu, Evgenii Zheltonozhskii, Terry Yue Zhuo, Thomas Wang, Olivier Dehaene, Mishig Davaadorj, Joel Lamy-Poirier, João Monteiro, Oleh Shliazhko, Nicolas Gontier, Nicholas Meade, Armel Zebaze, Ming-Ho Yee, Logesh Kumar Umapathi, Jian Zhu, Benjamin Lipkin, Muhtasham Oblokulov, Zhiruo Wang, Rudra Murthy, Jason Stillerman, Siva Sankalp Patel, Dmitry Abulkhanov, Marco Zocca, Manan Dey, Zhihan Zhang, Nour Fahmy, Urvashi Bhattacharyya, Wenhao Yu, Swayam Singh, Sasha Luccioni, Paulo Villegas, Maxim Kunakov, Fedor Zhdanov, Manuel Romero, Tony Lee, Nadav Timor, Jennifer Ding, Claire Schlesinger, Hailey Schoelkopf, Jan Ebert, Tri Dao, Mayank Mishra, Alex Gu, Jennifer Robinson, Carolyn Jane Anderson, Brendan Dolan-Gavitt, Danish Contractor, Siva Reddy, Daniel Fried, Dzmitry Bahdanau, Yacine Jernite, Carlos Muñoz Ferrandis, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. 2023. StarCoder: may the source be with you! <http://arxiv.org/abs/2305.06161>
- [25] Zhicheng Lin. 2024. How to write effective prompts for large language models. *Nature Human Behaviour* (2024), 1–5.
- [26] Anton Lozhkov, Raymond Li, Loubna Ben Allal, Federico Cassano, Joel Lamy-Poirier, Nouamane Tazi, Ao Tang, Dmytro Pykhtar, Jiawei Liu, Yuxiang Wei, et al. 2024. StarCoder 2 and The Stack v2: The Next Generation. *arXiv preprint arXiv:2402.19173* (2024).
- [27] Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng-Xin Yong, Hailey Schoelkopf, et al. 2022. Crosslingual generalization through multitask finetuning. *arXiv preprint arXiv:2211.01786* (2022).
- [28] Kawser Wazed Nafi, Tonny Shekha Kar, Banani Roy, Chanchal K. Roy, and Kevin A. Schneider. 2019. CLCDSA: Cross Language Code Clone Detection using Syntactical Features and API Documentation. In *2019 34th IEEE/ACM International*

- Conference on Automated Software Engineering (ASE)*. IEEE, San Diego, CA, USA, 1026–1037. doi:10.1109/ASE.2019.00099
- [29] Tasuku Nakagawa, Yoshiaki Higo, and Shinji Kusumoto. 2021. NIL: large-scale detection of large-variance clones. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. ACM, Athens Greece, 830–841. doi:10.1145/3468264.3468564
 - [30] Marianna Nezhurina, Lucia Cicolina-Kun, Mehdi Cherti, and Jenia Jitsev. 2024. Alice in Wonderland: Simple Tasks Showing Complete Reasoning Breakdown in State-Of-the-Art Large Language Models. *arXiv preprint arXiv:2406.02061* (2024).
 - [31] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research* 12 (2011), 2825–2830.
 - [32] Ruchir Puri, David S. Kung, Geert Janssen, Wei Zhang, Giacomo Domeniconi, Vladimir Zolotov, Julian Dolby, Jie Chen, Mihir Choudhury, Lindsey Decker, Veronika Thost, Luca Buratti, Saurabh Pujar, Shyam Ramji, Ulrich Finkler, Susan Malaika, and Frederick Reiss. 2021. CodeNet: A Large-Scale AI for Code Dataset for Learning a Diversity of Coding Tasks. <http://arxiv.org/abs/2105.12655>
 - [33] C.K. Roy and J.R. Cordy. 2008. NICAD: Accurate Detection of Near-Miss Intentional Clones Using Flexible Pretty-Printing and Code Normalization. In *2008 16th IEEE International Conference on Program Comprehension*. IEEE, Amsterdam, 172–181. doi:10.1109/ICPC.2008.41
 - [34] Chanchal Kumar Roy and James R Cordy. [n.d.]. A Survey on Software Clone Detection Research. ([n.d.]).
 - [35] Chanchal K Roy, James R Cordy, and Rainer Koschke. 2009. Comparison and evaluation of code clone detection techniques and tools: A qualitative approach. *Science of computer programming* 74, 7 (2009), 470–495.
 - [36] Pranab Sahoo, Ayush Kumar Singh, Sriparna Saha, Vinija Jain, Samrat Mondal, and Aman Chadha. 2024. A Systematic Survey of Prompt Engineering in Large Language Models: Techniques and Applications. *arXiv preprint arXiv:2402.07927* (2024).
 - [37] Vaibhav Saini, Farima Farmahinifarahani, Yadong Lu, Pierre Baldi, and Cristina V. Lopes. 2018. Oreos: detection of clones in the twilight zone. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. ACM, 354–365. doi:10.1145/3236024.3236026
 - [38] Hitesh Sajani, Vaibhav Saini, Jeffrey Svajlenko, Chanchal K. Roy, and Cristina V. Lopes. 2016. SourcererCC: scaling code clone detection to big-code. In *Proceedings of the 38th International Conference on Software Engineering*. ACM, Austin Texas, 1157–1168. doi:10.1145/2884781.2884877
 - [39] Nikita Sorokin, Dmitry Abulkhanov, Sergey Nikolenko, and Valentin Malykh. 2023. CCT-Code: Cross-Consistency Training for Multilingual Clone Detection and Code Search. <http://arxiv.org/abs/2305.11626>
 - [40] Chenning Tao, Qi Zhan, Xing Hu, and Xin Xia. 2022. C4: contrastive cross-language code clone detection. In *Proceedings of the 30th IEEE/ACM International Conference on Program Comprehension*. ACM, Virtual Event, 413–424. doi:10.1145/3524610.3527911
 - [41] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madsen Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. <http://arxiv.org/abs/2307.09288>
 - [42] Lewis Tunstall, Nathan Lambert, Nazneen Rajani, Edward Beeching, Teven Le Scao, Leandro von Werra, Sheon Han, Philipp Schmid, and Alexander Rush. 2023. Creating a Coding Assistant with StarCoder. *Hugging Face Blog* (2023).
 - [43] Tijana Vislavski, Gordana Rakic, Nicolas Cardozo, and Zoran Budimac. 2018. LICCA: A tool for cross-language clone detection. In *2018 IEEE 25th International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, Campobasso, 512–516. doi:10.1109/SANER.2018.8330250
 - [44] Wenhan Wang, Ge Li, Bo Ma, Xin Xia, and Zhi Jin. 2020. Detecting Code Clones with Graph Neural Network and Flow-Augmented Abstract Syntax Tree. In *2020 IEEE 27th International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, London, ON, Canada, 261–271. doi:10.1109/SANER48275.2020.9054857
 - [45] Huihui Wei and Ming Li. 2017. Supervised Deep Features for Software Functional Clone Detection by Exploiting Lexical and Syntactical Information in Source Code. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, Melbourne, Australia, 3034–3040. doi:10.24963/ijcai.2017/423

- [46] Yuxiang Wei, Federico Cassano, Jiawei Liu, Yifeng Ding, Naman Jain, Harm de Vries, Leandro von Werra, Arjun Guha, and Lingming Zhang. 2024. StarCoder2-Instruct: Fully Transparent and Permissive Self-Alignment for Code Generation. <https://github.com/bigcode-project/starcoder2-self-align>
- [47] Martin White, Michele Tufano, Christopher Vendome, and Denys Poshyvanyk. 2016. Deep learning code fragments for code clone detection. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering*. ACM, Singapore Singapore, 87–98. doi:10.1145/2970276.2970326
- [48] Qinyuan Ye, Maxamed Axmed, Reid Pryzant, and Fereshte Khani. 2024. Prompt Engineering a Prompt Engineer. <http://arxiv.org/abs/2311.05661>
- [49] Ziqi Yin, Hao Wang, Kaito Horio, Daisuke Kawahara, and Satoshi Sekine. 2024. Should We Respect LLMs? A Cross-Lingual Study on the Influence of Prompt Politeness on LLM Performance. *arXiv preprint arXiv:2402.14531* (2024).
- [50] Gang Zhao and Jeff Huang. 2018. DeepSim: deep learning code functional similarity. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. ACM, 141–151. doi:10.1145/3236024.3236068
- [51] Ming Zhu, Aneesh Jain, Karthik Suresh, Roshan Ravindran, Sindhu Tipirneni, and Chandan K. Reddy. 2022. XLCoST: A Benchmark Dataset for Cross-lingual Code Intelligence. <http://arxiv.org/abs/2206.08474>

Received 2024-09-13; accepted 2025-01-14