

REINFORCEST: Reinforcing Semantic Code Similarity for Cross-Lingual Code Search Models

Anthony Saieva[†]

IBM Research

New York, NY, USA

anthony.saieva@ibm.com

Saikat Chakraborty[†]

Microsoft Research

Redmond, WA, USA

saikatc@microsoft.com

Gail Kaiser

Columbia University

New York, NY, USA

kaiser@cs.columbia.edu

Abstract—This paper introduces a novel code-to-code search technique that enhances the performance of Large Language Models (LLMs) by including both static and dynamic features as well as utilizing both similar and dissimilar examples during training. We present the first-ever code search method that encodes dynamic runtime information during training without the need to execute either the corpus under search or the search query at inference time and the first code search technique that trains on both positive and negative reference samples. To validate the efficacy of our approach, we perform a set of studies demonstrating the capability of enhanced LLMs to perform cross-language code-to-code search.

Our evaluation demonstrates that the effectiveness of our approach is consistent across various model architectures and programming languages. We outperform the state-of-the-art cross-language search tool by up to 44.7%. Moreover, our ablation studies reveal that even a single positive and negative reference sample in the training process results in substantial performance improvements demonstrating both similar and dissimilar references are important parts of code search. Importantly, we show that enhanced well-crafted, fine-tuned models consistently outperform enhanced larger modern LLMs without fine tuning, even when enhancing the largest available LLMs highlighting the importance for open-sourced models.

To ensure the reproducibility and extensibility of our research, we present an open-sourced implementation of our tool and training procedures called REINFORCEST.

Index Terms—Large Language Models, Semantic Clones, Semantic Code Search

I. INTRODUCTION

With the rise of powerful large language models (LLMs) since the original BERT [1] paper was published, machine learning applied to source code analysis tasks has become increasingly popular in the last few years. At the same time as large public repositories of code have become available [2], the potential of code-to-code search has gained new significance to aid in common software maintenance and development tasks like code migration [3], transpilation [4], code repair [5], bug detection [6], education [7], and refactoring [8].

In this paper, we present a series of novel techniques that enhance LLMs for code-to-code search and related tasks. We introduce a new code-to-code search technique that encodes both source code and dynamic runtime information during training without needing to execute anything for inference from either the search corpus or the search query. Figure 1

shows the general premise. First, during training the model learns both static and dynamic similarity from the training corpus, maximizing the distance between dissimilar code and minimizing the distance between similar code. We generate a single dynamic similarity feature we call the Semantic Similarity Score (SSS) before training that involves executing as much of the training corpus as possible with the same inputs and comparing the outputs of functions that have the same input types. This is the only time code is executed. Then to put this model to use, the model is used to create a precomputed embedding of the search corpus, and finally at inference time the model embeds the search query so a simple comparison to the precomputed embeddings yields the search results. Traditionally dynamic analysis techniques require executing code during training as well as executing the search corpus and search query during the search procedure. This leads to problems scaling and all of the environment-related issues that come with running code in practice [9]. In contrast our approach only incurs the overhead of executing training set code in a controlled environment, and once the model is distributed code never needs to be executed again. Notably, our approach still accommodates circumstances when parts or all of the training set are not executable.

To validate the effectiveness of our approach, we considered the problem of cross-language search in which the query and corpus are from different languages, as in COSAL [10]. We felt this would best demonstrate the effectiveness of our technique since we are most concerned with behavioral similarity and searching across syntactically different languages means the model cannot rely on source code similarity.

We conduct a series of ablation studies that demonstrate the ability of LLMs to infer dynamic behavior under a variety of circumstances even when inference is based solely on static information. These circumstances include multiple query and corpus languages, various underlying LLM models both when the models are open-sourced and proprietary, when dynamic runtime information is available and without runtime information, and finally varying the number of both positive and negative reference samples in training. Our evaluation shows that our technique performs considerably better than state-of-the-art cross-language search techniques by up to 44.7% on the benchmark Atcoder [11] dataset, which consists of Java and Python solutions corresponding to 361 problems

[†]Equal Contribution

from 83 programming contests, and maintains near or better performance with a variety of LLM architectures including CodeBERT [12] and OpenAI’s codex models [13], even when the LLM’s cannot be fine-tuned. It is worth noting that we find that a smaller fine-tuned LLM outperforms significantly larger proprietary LLM’s highlighting the value and importance of open-sourced models for the research community. Further, we find that as long as there is at least one positive reference and one negative reference available during training our technique performs effectively compared to benchmarks, allowing for sparser datasets as what might be found in practice [14] and demonstrating the importance of considering both similar and dissimilar reference samples.

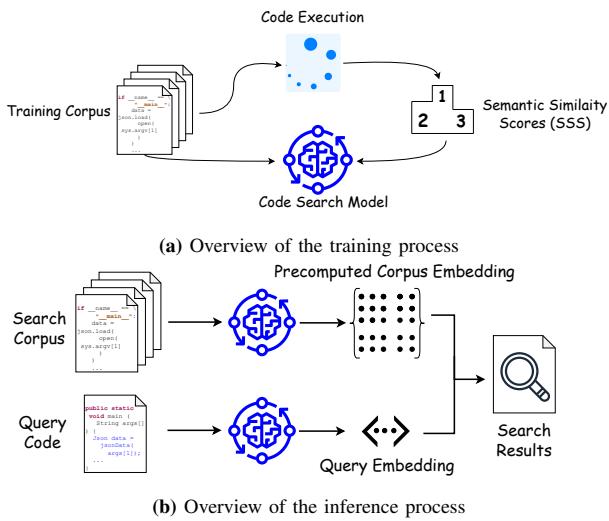


Fig. 1: Overview of the training and inference processes

We make the following contributions:

- 1) The first code-to-code search technique that encodes dynamic runtime information during training without the need to execute any code from either the search corpus or the search query during inference.
- 2) The first code-to-code search technique that considers both similar and dissimilar reference samples during training.
- 3) A set of ablation studies that demonstrate the ability of LLM’s to infer dynamic behavior during cross-language search even though inference is only based on static information.
- 4) An evaluation that shows our technique remains effective across many model architectures.
- 5) Ablation studies showing even a single positive and single negative reference sample in training results in effective performance.
- 6) An evaluation showing well-crafted fine-tuned models out-perform larger non fine-tuned LLM models even with the largest LLM’s.
- 7) An open-sourced implementation of our tool and training

procedures called REINFOREST ¹ that represents a new state-of-the-art in cross-language code search.

II. BACKGROUND

A. Code Search

Code-to-code search can be defined broadly as taking a code sample (the query), and searching through a corpus of other available code samples (the search corpus), and finding code clones in the search corpus matching the query. Cross-language code-to-code search deals specifically with situations when the query and corpus are written in different languages. Consider a situation in which Jane the developer is tasked with writing code in a language she is relatively new to or a language that lacks supporting documentation. Rather than writing without a guide, with cross-language code-to-code search Jane can make a query with a reference sample from a language she knows well and look for similar samples in a corpus of the language she is not as comfortable with.

Su et al. [15] define four type of code clones:

- **Type 1:** Identical code fragments, except for variations in whitespace, comments, or formatting.
- **Type 2:** Syntactically identical code fragments, with differences in variable, constant, or function names.
- **Type 3:** Similar code fragments with modified statements, function calls, or control structures.
- **Type 4:** Semantically identical code fragments with the same functionality but different syntactic structures.

Traditionally, type four clones are the most difficult to detect as the source code may offer little if any useful context, and that is the type we are most concerned with in this work. If our technique successfully learns code behavior we should be able to perform code search even when the syntactic structures are from different languages.

Code clones arise from several factors, including copy-pasting, code reuse, and adherence to coding patterns, and code search techniques are essential tools for developers, as they help locate and reuse code snippets in large codebases. Existing static code search techniques can broadly be classified into three categories 1) text-based approaches: These methods rely on traditional information retrieval techniques such as keyword search or string matching to find relevant code snippets, 2) structural approaches: These methods leverage the syntactic and semantic structure of source code, such as abstract syntax trees, to find matching or similar code fragments, and 3) machine learning-based approaches, which leverage both text and semantic structural encoding along with machine learning algorithms to find clones. Traditionally code search has been limited to query and corpus of code from the same language, but in recent years this has expanded to search code bases where the query and corpus may be of different languages [16]–[18] as code search algorithms get better at distilling semantic information from the source code text.

Additionally techniques that include information from dynamic analysis like execution traces [19] or input output pairs

¹<https://github.com/reinforest-team/REINFOREST>

[15] leverage additional context to find code clones that were undetected by static analysis. This is especially important in type four clones, which only concerns behavioral similarity.

On the one hand static code search techniques do not require code execution, making them highly scalable for searching large codebases processing vast amounts of source code quickly, and providing fast search results to developers [20]. They are also easy to implement as text-based and structural approaches in static code search are generally simpler to implement than dynamic methods, and require no execution environment nor runtime setup making them more portable and applicable to a wide range of software projects. These practical benefits come at the cost of less context than what is available to dynamic techniques, which results in lower performance. On the other hand, dynamic techniques offer improved accuracy and context-awareness but can be more complex and less scalable [21].

By including dynamic information during training where possible, and not requiring it at inference time, we aim to bridge the gap between dynamic and static code search bringing together both the convenience of static methods and the enhanced context of dynamic methods.

B. Encoder Models for source code understanding

With the advancement of Deep Learning models and techniques, models have been trained to "understand" source code. Given a code snippet c and an encoder Enc , the model generates a d dimensional vector representation of the code $Enc(c) \in \mathbb{R}^d$. This vector represents the lexicographic, syntactic, and semantic information in the code. Such vector embedding can be subsequently used for various different tasks including, but not limited to, classification, search, ranking, etc. The encoder can be an open-source model, such as CodeBERT [12] or a closed-source embedding API such as OpenAI's codex embedding [22]. For a highly effective pretrained model, the embeddings themselves should be sufficient for tasks related to code understanding, however as previous research has shown [23], the embeddings need to be tuned depending on the downstream tasks. In the case of open-source models for embedding, where we can access the models' weights, the models can be fine-tuned as part of the downstream tasks. For embeddings coming from closed-source models, further models can be built on top of the initial embeddings, but the embeddings cannot be changed.

III. REFOREST

In this section, we describe the code search technique we propose in REFOREST. We describe REFOREST as two disjoint steps – (a) Training and (b) Online query embedding and search. At the code level, REFOREST consists of two different encoders – query code encoder (E_q) and the document code encoder (E_d). Each of these, given an input code x , generates a fixed-sized vector representation of x .

A. Training REFOREST

We train REFOREST aiming to optimize for two different orthogonal goals. First, we would want REFOREST to

transform embeddings generated by the query code encoder (E_q) and the document code encoder (E_d) in such a way that minimizes the distance between related code and maximizes the distance between unrelated code. Our second goal is to endow REFOREST with dynamic code similarity information during training so that, during inference (search), we do not have to calculate the semantic (I/O) similarity between the query code and all other code in the database. Figure 2 shows an overview of REFOREST's training process.

Contrastive Training. To achieve the first goal, we take inspiration from contrastive learning [24]. We assume the presence of a dataset \mathcal{D} consisting of a set of tuples $(c_s, \mathcal{C}_p, \mathcal{C}_n)$. c_s is a piece of code written in the source language s (e.g., Java). \mathcal{C}_p and \mathcal{C}_n are the set of code written in the target language (e.g., Python). The set $\mathcal{C}_p = \{p_{t1}, p_{t2}, \dots, p_{tk}\}$ are the solutions to the same problem as c_s . In contrast, $\mathcal{C}_n = \{n_{t1}, n_{t2}, \dots, n_{tk}\}$ are solutions to a different problem than c_s . We refer to the code samples \mathcal{C}_p as positive samples and code samples \mathcal{C}_n as negative samples hereafter. Intuitively, we would want to bring the embeddings of each element in \mathcal{C}_p closer to the embedding of c_s , and we would want to push the embedding of \mathcal{C}_n 's elements far away from c_s 's embedding.

First, we encode c_s with the query code encoder, E_q , and both the \mathcal{C}_p and \mathcal{C}_n sets with the document code encoder E_d . Formally, $R_s = E_q(c_s)$, $\mathcal{R}_p = \{E_d(p_{ti}) \mid p_{ti} \in \mathcal{C}_p\}$, and $\mathcal{R}_n = \{E_d(n_{ti}) \mid n_{ti} \in \mathcal{C}_n\}$, where R_s is the embedding of the query code, \mathcal{R}_p , and \mathcal{R}_n are the sets of embedding corresponding to positive and negative samples, respectively. Then we compute the cosine similarity between the query code and each of the positive and negative samples. We compute the cosine similarity between two vectors, a and b as,

$$sim(a, b) = \frac{|a \cdot b|}{|a| * |b|} \quad (1)$$

Our next goal is to simultaneously maximize the similarity between c_s and all elements in \mathcal{C}_p and minimize the similarity between c_s and all elements in \mathcal{C}_n . We achieve such an objective by minimizing the following loss function,

$$\mathcal{L} = \sum_{(c_s, \mathcal{C}_p, \mathcal{C}_n) \in \mathcal{D}} \left(\sum_{r_{ti}^p \in \mathcal{R}_p} (l_p - sim(R_s, r_{ti}^p))^2 \right) + \left(\sum_{r_{ti}^n \in \mathcal{R}_n} (l_n - sim(R_s, r_{ti}^n))^2 \right) \quad (2)$$

Here, l_p and l_n are target labels for positive and negative samples. The values of l_p and l_n should depend on the similarity function being used in the system. For an unbounded similarity function, l_p and l_n should be set to ∞ and $-\infty$, respectively. In the case of REFOREST, we set $l_p = 1$ and $l_n = 0$, since these values are the maxima and minima of the similarity function we use (Equation (1)).

Semantic Similarity Score As our second training goal is to endow the models with the knowledge about semantic (I/O) similarity, we assume a function $\mathcal{S}_{io}(c_s, c_t)$, which computes the semantic similarity between arbitrary code pieces c_s and

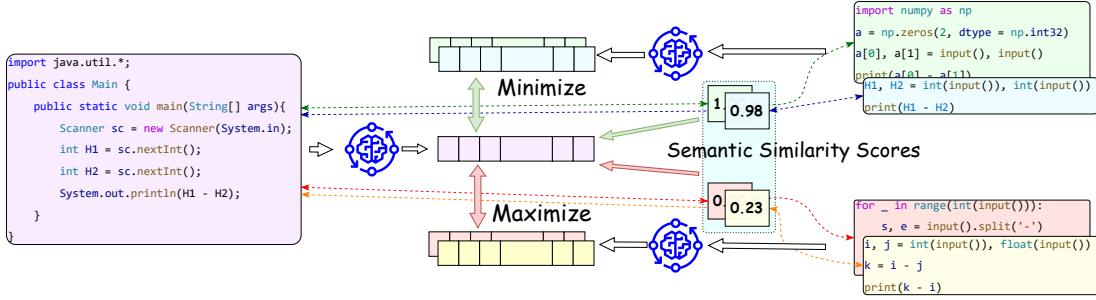


Fig. 2: Training process in REINFORCEST

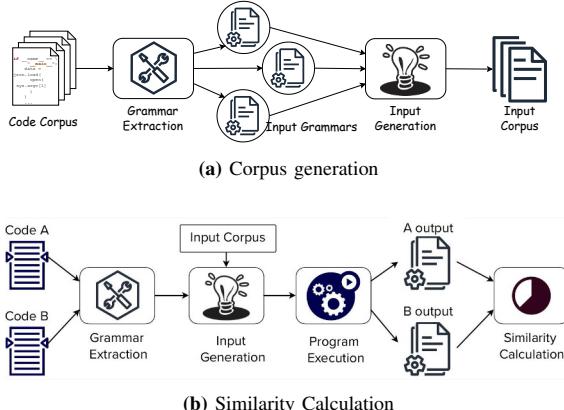


Fig. 3: Similarity Score Processes

c_t . We endow the model with such semantic similarity by modifying the loss function in Equation (2) as

$$\mathcal{L}' = \sum_{(c_s, c_p, c_n) \in \mathcal{D}} \left(\sum_{p_{ti}, r_{ti}^p} (l'_p(c_s, p_{ti}) - \text{sim}(R_s, r_{ti}^p))^2 \right) + \left(\sum_{n_{ti}, r_{ti}^n} (l'_n(c_s, n_{ti}) - \text{sim}(R_s, r_{ti}^n))^2 \right) \quad (3)$$

$$l'_p(c_s, p_{ti}) = (1 - \alpha)l_p + \alpha S_{io}(c_s, p_{ti}) \quad (4)$$

$$l'_n(c_s, n_{ti}) = (1 - \alpha)l_n + \alpha S_{io}(c_s, n_{ti}) \quad (5)$$

Here, α represents the relative importance of semantic training. In an ideal scenario, where we could compute the semantic similarity between all pairs of c_s and c_t , we argue that setting α to 1.0 would result in the best performance. However, it is practically infeasible to assume that for all such pairs of c_s and c_t , $S_{io}(c_s, c_t)$ exists. Thus, we empirically found best validation performance by setting α to 0.2. Regardless, α remains a user-defined hyperparameter in REINFORCEST's implementation.

Semantic Similarity Score Calculation For Training. REINFORCEST's architecture and process uses a semantic similarity score calculation based on runtime information where possible during training. We use SLACC's [25] similarity

scoring methodology, which is simply the number of matching outputs divided by the number of inputs. Figure 3 outlines our two process procedures for generating similarity scores. First, as shown in Figure 3a, we generate the input corpus from the code corpus. For any piece of code, we statically analyze the system calls and derive the structure of the input for that particular piece of code. After extracting all of the input structures in the corpus, we generate a large number of inputs of random values for each structure based on primitive types and create an input corpus. Then as shown in 3b for any two samples, we extract the input structures, run them with the inputs that match their input structure, and then perform the similarity calculation described earlier. As in SLACC, if two pieces of code have different input structures, they receive a similarity score of zero.

B. Code Search With REINFORCEST

Once the training process finishes, we have two trained encoders, E_q and E_d , to encode queries and documents, respectively. Immediately after the training, we encode all the code in the search database with E_d . Since the encoded representations are real-valued vectors, indexing tools such as FAISS [26] can store them in the indexed database for efficient search. For every query code c_s , we generate the query embedding R_s with the query encoder E_s . Then we search in the embedding database based on the similarity score in Equation (1) and return the top n candidate with the highest similarity scores.

IV. EXPERIMENTAL DESIGN

A. Research Questions

To evaluate our technique we implemented it in a tool called REINFORCEST and we considered four main questions. The most important and broadest of which is simply:

RQ1. How does REINFORCEST's performance compare to the performance of other cross-language code search techniques?

Still better performance could mean that we simply rely on the power of modern LLM's. If, however, our tool continues to show state-of-the-art performance with multiple LLM's, and our additional training method always improves performance this demonstrates that our technique is in fact the source of any performance advantages shown in RQ1. This leads to a second question:

TABLE I: Statistics of the dataset. The “Average” denotes the average number of files per problem.

	Train	Valid	Test
Number of Problems	287	37	37
Number of Java files	Total	14413	2533
	Average	50.22	68.50
Number of Python files	Total	17182	2983
	Average	59.87	80.62
			58.16

RQ2. Does REINFORCEST’s methodology and performance generalize across different models?

Even if RQ2 shows performance improvements across multiple architectures, it may only be from source code analysis and the dynamic runtime information may not make a difference, so we consider:

RQ3. Does including semantic similarity scores during training improve code search?

Lastly, it remains unclear what impact including both positive and negative samples during training has, so lastly we evaluate:

RQ4. How does changing the number of positive and negative comparison samples available for training effect REINFORCEST’s performance?

B. Dataset

To evaluate REINFORCEST, we used the Atcoder dataset [11], which consists of 18644 Java solutions and 22317 Python solutions corresponding to 361 programming contest problems. To prove our concept, we created the train, validation, and test splits by dividing the dataset across different problems – *i.e.*, train, valid, and test splits do not share the solution from the same problem. Since the problems are independent of each other, dividing in such a way prohibits data leakage across splits. Across all the experiments presented in the paper, we evaluate based on the same splits. Table I shows detailed statistics of the dataset. All experiments were performed with an A100 NVIDIA with 216 GB of RAM with a 2.44 GHz AMD 64-core processor.

C. Evaluation Metrics

We used three different evaluation metrics. The first is precision at N (PR@N), where N is between one and five inclusively. Precision at N is calculated as the average number of truly positive examples in the first N search results using every sample in the test set as a query against the rest of the test set. These metrics are benchmark standards in code search [27], but we feel they are most important as they are indicative of the experience a user would have looking at a list of search results following a query. We also examined a metric termed Average Rank Gap (ARG) [28], which is defined as the average rank of negative examples minus the average rank of positive examples. Unlike PR@N metrics, the ARG gives a sense of the entire set of search results available in the corpus. Between the PR@N metric and the ARG metric we get a sense of what a user would experience as well as REINFORCEST’s

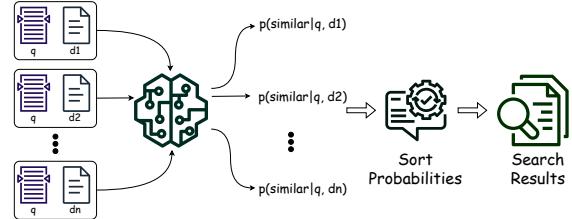


Fig. 4: Classifier Style Search

performance across the entire dataset. Lastly we evaluated the tools based on the average first position (AFP) [29] of a positive result (lower is better). ARG provides context for the entire result set while AFP provides context for the highest ranked portion of the result set as a whole.

V. EMPIRICAL RESULTS

We answer the questions in the previous section one by one here. A full table of our results has been omitted for space consideration and clarity but is included in the supplementary material.

A. RQ1 - Overall Performance

First we examined the overall performance of the model, *i.e.*, how good is the model at finding cross-language clones. Using the evaluation metrics defined in the previous section, we evaluated against both techniques that require training as well as those that do not. For strawman comparison we implemented a TF-IDF search with the BM25 python library [30] that treats source code like natural language and performs a token-based search. We also implemented an AST-based search adapted from SLACC [31], in which we take an AST representation of both the query and the document in the corpus, and convert them from language-specific AST’s to a generic AST with common node types where possible. Then we use the standard tree distance algorithm [32] to determine the similarity of the two samples. In the last of our non-training comparisons, we reimplemented COSAL’s [10] token subset technique in which functions are divided into smaller snippets and the smaller snippets are clustered to determine a similarity between functions. We reimplemented this system because it is the state-of-the-art in cross-language code search, and we have different train, validate, test splits on the Atcoder dataset compared to their Atcoder experiments and we have a different experimental setup.

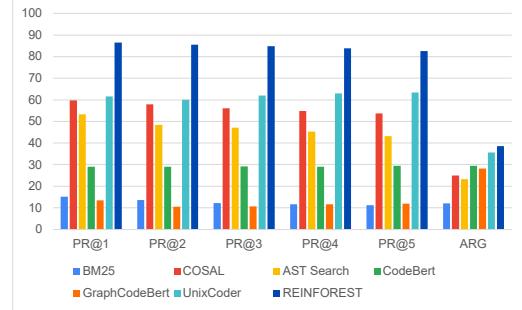
To compare against other code search techniques that require training, we selected three state-of-the-art LLM models that have shown promising results in a variety of code related tasks. CodeBERT [12] is a pre-trained model for code representation and generation developed by Microsoft Research Asia. It can be fine-tuned for various downstream tasks and has achieved state-of-the-art performance on several benchmark datasets. CodeBERT’s ability to transfer knowledge across programming languages makes it particularly useful for code generation tasks and has many potential applications in software engineering and programming language

research. GraphCodeBERT [12] is an extension of CodeBERT that incorporates control flow graphs (CFGs) into its pre-training process. By using CFGs, GraphCodeBERT can better capture the relationships between code tokens and produce more accurate representations for code analysis and generation tasks. It has achieved state-of-the-art performance on several code analysis benchmarks, demonstrating its effectiveness in modeling complex code structures. UnixCoder [33] is a pre-trained model for natural language to shell command translation, developed by Facebook AI. It is trained on a large corpus of command-line usage examples and can generate shell commands given a natural language input. UnixCoder has achieved state-of-the-art performance on several benchmarks, demonstrating its ability to understand and translate complex natural language commands into executable shell commands.

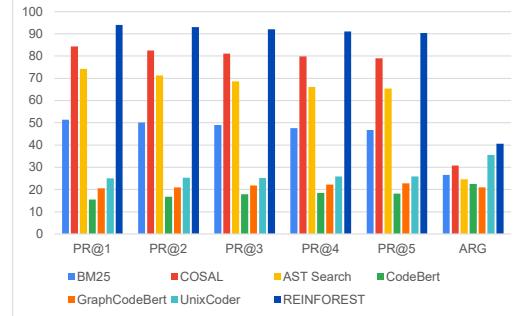
We implemented a search algorithm with each of the three models as described in Figure 4, similar to that of [34]. The models are trained as a classifier on the training set and then for every document in the test set the classifier determines the probability of the document matching the given query and sorts results based on the highest probability. As such at inference time this process requires classifying the entire test corpus and incurs tremendous computational overhead. REINFOREST in contrast can use a precomputed corpus embedding as shown in Figure 1b and only needs to create an embedding vector from the query without concerning the rest of the corpus. To get the search results, simply choose the precomputed embeddings that are closest to the query vector. Performance experiments for our best version of REINFOREST took roughly 14 hours while the other trained model code search comparisons would have taken multiple days to complete. To deal with this we randomly sampled 200 queries from the test set and reported our findings for those 200 samples. With this stratified sampling the experiments took roughly the same amount of time as the experiments with REINFOREST.

Figure 5 details the results of our experiments for PR@N metrics and ARG with the best-performing version of REINFOREST against all the comparison techniques. For visual clarity, PR@N metrics were normalized to a percentage scale, ARG metrics were scaled by 100x, and AFP metrics were omitted but original values are included in the supplemental material. Figure 5a shows results for a Java query against a corpus of Python code while Figure 5b shows results for a Python query against a corpus of Java code. Across all evaluation metrics REINFOREST outperforms all comparison techniques with improved performance up to 40% compared to the next best performing method. Somewhat surprisingly we see the no training required techniques outperforming the trained techniques with the exception of UnixCoder and Java to Python search. While the relative performance of each method changes depending on the evaluation criteria, we find unsurprisingly that current state of the art COSAL generally is the next best performing.

Since REINFOREST outperforms existing cross-language code search techniques by a considerable margin on all evaluation metrics we determine the answer to RQ1 as:



(a) Java to Python Search



(b) Python to Java Search

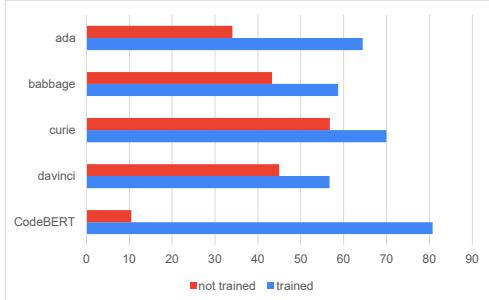
Fig. 5: Overall REINFOREST Performance

Result 1: REINFOREST performs considerably better than existing techniques with a 40.6% improvement over the nearest baseline for Java to Python search, and an 11.5% improvement over the nearest baseline in Python to Java search. We outperform COSAL’s state-of-the-art performance by up to 44.7%.

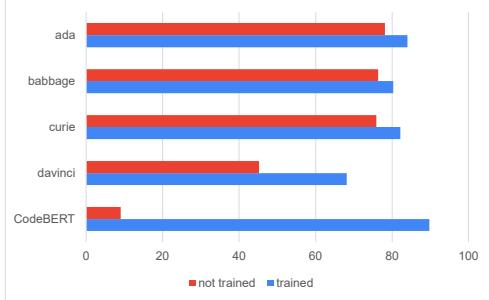
B. RQ2 – Model Generality

While the results from the previous section show that the techniques described result in an effective tool, it remains unclear if the performance is because of the training techniques we employed or just the power of modern LLM’s. To examine this we ask in RQ2 if our REINFOREST techniques are effective in improving performance across many different LLM models both open-sourced and proprietary.

For performance comparison we use the same metrics as described in the last subsection: PR@N, ARG, and AFP. We chose the most advanced available LLM models for comparison. First CodeBERT for the reasons described previously, but instead of training it as a classifier as we did in the first RQ, we fine-tuned its embedding on the cross-language code search task during the training procedure described in Section IV. Additionally we ran our training models on OpenAI’s Codex Ada, Babbage, Curie, and Davinci models. These are proprietary models and as such we could not perform fine-tuning during training, so we used the encodings from these models as is during training and REINFOREST learns to maximize and minimize distances in these encoding spaces. Without access to the model internals the only discernible difference between the OpenAI models is the vector size of the embeddings, which are Ada: 1024, Babbage: 2048, Curie:



(a) Java Training Improvement (PR@1)



(b) Python Training Improvement (PR@1)

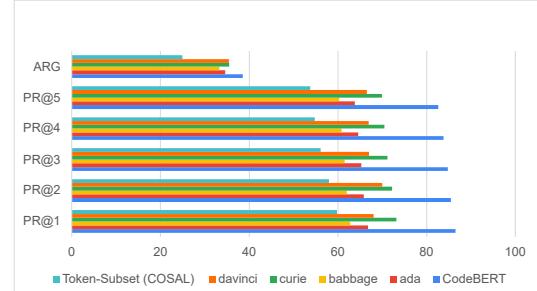
Fig. 6: REINFORCEST Training Impact (PR@1)

4096, and Davinci: 12188. As such we cannot explain why some OpenAI models may perform better than others.

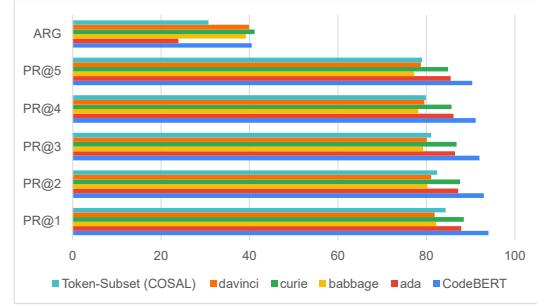
To answer this question we ran two experiments. First we compared the performance of the various models' embeddings on cross-language code-to-code search without any alterations to REINFORCEST's performance on the same task training on the embeddings from those same models. Any increase in performance would have to be attributed to REINFORCEST's procedures instead of the base LLM.

Second we compared the performance of REINFORCEST trained with the various underlying model embeddings on each evaluation metric to the next best performing code-to-code search technique. To understand the significance of the generality of our technique we should examine it relative to the state-of-the-art across all model architectures. Achieving state-of-the-art performance with all underlying models would further suggest that our techniques are valuable independent of the underlying LLM's.

Figure 6 compares performance between trained and untrained models. Clearly REINFORCEST's training always improves code-to-code search performance for both Java to Python and Python to Java queries. The improvements range from 5.2% all the way up to 1044%. Unsurprisingly our training has the largest impact on the one model that we could fine-tune during the training procedure. More interestingly, however, we find that despite having the weakest performance without training the fine-tuned CodeBERT model outperforms all other models on both Java and Python cross-language queries. This shows that fine-tuning the embedding during the training procedure has a larger impact than the size of the



(a) Java Model Performance



(b) Python Model Performance

Fig. 7: Performance of REINFORCEST across different models

overall LLM and demonstrates the importance of fine-tuning.

Figure 7 shows the overall performance of the various models using REINFORCEST's techniques relative to COSAL (SOTA). Figure 7a details the results of Java to Python queries and shows that REINFORCEST achieves SoA performance across all metrics. Interestingly the same cannot be said for Python to Java queries as COSAL outperforms REINFORCEST in some circumstances but never outperforms all versions of REINFORCEST. Since we do not know the internals of OpenAI's models it remains unclear why some of their models perform better than the SoA while others do not. However, most importantly we find that fine-tuned CodeBERT outperforms all other models on every metric for both cross-language queries. This solidifies fine-tuned CodeBERT as the highest performing version of REINFORCEST and demonstrates the power of open-sourced models. It also suggests that if the proprietary models were open-sourced we might see greater improvement.

The results from these experiments show 1) REINFORCEST's training techniques always improve the base LLM models and 2) REINFORCEST's techniques outperform the state-of-the-art in most circumstances and always when fine-tuning is possible. Thus we find that:

Result 2: REINFORCEST's techniques generalize across multiple LLM models, improving performance by a minimum of 26.2% on Java to Python search and 5.17% on Python to Java search, and when the base LLM is fine-tuned during training performance improves by 7.69x on Java to Python search and 9.96x on Python to Java search.



(a) Java Impact of SSS (PR@1)



(b) Python Impact of SSS (PR@1)

Fig. 8: REINFORCEST SSS Impact (PR@1)

C. RQ3 – Impact of SSS

The previous RQ results show that REINFORCEST’s techniques are in fact responsible for the performance in RQ1, we don’t yet know the impact of including runtime data in the form of the semantic similarity score (SSS). The SSS is the scored input output comparison described in Section IV.

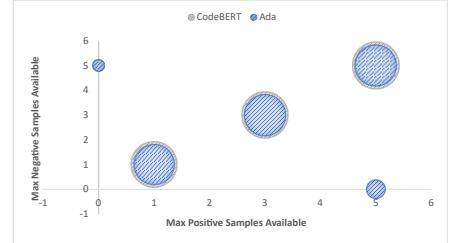
We evaluated the SSS’s impact by comparing REINFORCEST with and without SSS in training. Any performance gain would indicate the model learned runtime behavior from SSS. We tested this with CodeBERT and OpenAI models to determine if learning runtime behavior is a general trait of modern LLMs or specific to one model.

Figure 8 shows that including SSS consistently boosts REINFORCEST’s performance across models and languages. This improvement, seen even without code execution during inference, indicates that learning dynamic behavior is a general trait of modern LLMs, not limited to a single model. The most significant gain was with OpenAI’s davinci model, suggesting this technique enhances performance regardless of fine-tuning.

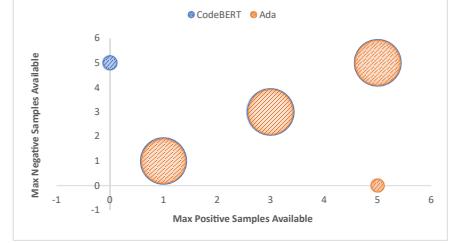
Result 3: Including SSS improves REINFORCEST’s code search performance across all embeddings. Including the SSS with the best performing version of REINFORCEST (CodeBERT) contributed to an average of 7% improvement for Java to Python queries and a 4.8% improvement for Python to Java queries.

D. RQ4 – Varying Training Samples

Previous experiments show that 1) REINFORCEST’s overall performance is better than the state-of-the-art, 2) is consistent regardless of the underlying model architecture, and 3) leverages runtime behavioral features. However it remains unclear if the model is learning from both positive and negative



(a) Java to Python Search varying reference samples (PR@1)



(b) Python to Java search varying reference samples (PR@1)

Fig. 9: REINFORCEST performance with varying number of reference samples during training

reference samples during training or if it is simply relying on one or the other. Further we have not yet investigated the impact of each reference sample on the model’s performance. It seems plausible that increasing the size of the training set or adding more samples could improve performance.

So we investigate RQ4 by varying the maximum number of positive and negative samples available during training. If the model is only learning from positive samples then performance should not dip when negative samples are removed and vice versa. Further if more training data would improve performance then if the number of reference samples available during training increases performance should increase as well. We ran performance experiments with the same number of positive and negative reference samples for one, three, and five references. Then we ran the experiments with zero positive references and five negative references and vice versa. we did this with a CodeBERT model fine-tuned during training and all available OpenAI Codex models off the shelf as in RQ2.

Full results of these experiments are included in the supplementary materials with all models and evaluation metrics. Figure 9 shows these results with both Java to Python search and Python to Java search for CodeBERT and OpenAI’s Ada model evaluated on the PR@1 metric. The x and y axes show the number of positive and negative samples available during training and the size of the circles represents the PR@1 performance of the model in the circumstances described by the x and y axis. Larger is better.

There are two interesting findings consistent in both models that can be fine-tuned and those that can not. The first is that when either positive or negative samples are omitted during training, performance decreases drastically. This indicates that REINFORCEST’s dual maximizing and minimizing technique during training is in fact responsible for its performance. If

only the positive samples were responsible for performance, then removing the negative samples would have no impact. This is clearly not the case. Second, we see that performance doesn't vary greatly when a single positive and negative reference sample are available. This means that increasing the size of the training set would not have meaningful impact on results. Still we do see a slight increase indicating it has some impact. This could be explained by the fact that during training source code labels are binary, which means that additional positive and negative samples do not explicitly provide extra context. With a different training metric that quantified the distance between samples more explicitly it is possible that the number of samples available would have more impact. However, this finding indicates that even in the case of a sparse dataset available during training, REINFORCEST's performance is still effective even if only a single positive and negative reference sample is available.

Result 4: REINFORCEST's performance depends on both positive and negative samples. Using a combination of positive and negative samples improves performance by up to 10.2x and 17.8x for Python to Java search and 15.5x and 12.2x for Java to Python search, compared to only positive and negative references respectively. The addition of more than one positive and negative sample during training performance improvement decays, with a 5.4% and 0.6% improvement for Java to Python and Python to Java search, respectively.

VI. RELATED WORK

A. Code Search

Code search has been studied for a long time but has become increasingly more important with the rise of machine learning and the availability of large open-source code repositories like Github [35]. Di Grazia et al. provide a comprehensive overview [36]. Techniques have generally fallen into two groups. The first relies primarily on static techniques that do not run the code either at training or inference time and the second also includes dynamic information for more context but takes on the additional overhead of needing to run code both in training and at inference time. Early static code similarity solutions like CCFinder [37] and CP-Miner [38] rely on token-based intermediate representations, and other early solutions like [39] rely on graph-based and AST-based intermediate representations. Recently, FaCoy [40] uses code fragments to augment the code query to increase useful matches but encodes no behavioral information and requires previous queries to gain context. Pinku et al. show that additional information from a transcompiler can be powerful enough to take single language models and apply them across languages [41]. Adding such additional information remains cumbersome at inference time. Other tools focus on the ability to match structurally different but semantically similar code clones like CCGraph [42], which leverages program dependency graphs and graph kernels to run machine learning directly on a graph representation of the code. These inherently expensive representations are difficult to train and do not encode any dynamic information. Another tool SourcererCC [43] focuses on such matches in large code

bases, and while it scales well it fails to capture enough context to perform as well as other tools. Some security-focused code search tools like Macneto [44] apply clone detection techniques to an intermediate representation, e.g., examining Java byte code to detect clones after obfuscation. Such tools fall into a different category of clone detection outside of source code clone detection.

The other set of code clone detection tools leverage dynamic behavior in their code representations. While this provides important semantic information, it adds complexity to the problem by requiring runnable code both during training and at inference time. With the modern intricacies of build environments this is not practical at scale. Some recent work on dynamic code clone detection has moved towards taking code and splitting it into smaller segments and comparing the segments. This technique attempts to mitigate the difficulties of trying to match large code segments that are likely to behave similarly but have different structures, like COSAL [10] and SLACC [25]. The technique is especially valuable in comparing code across languages. But splitting the code struggles to scale as the corpus grows larger, especially in a dynamic context that executes fragments repeatedly for behavioral information. HitoshiIO [15] examines function input and output pairs to identify similar code and DyCLink [19] combines this information with code graph analysis and execution path information for a more informed but heavyweight solution. To date, as far as we know, REINFORCEST is the first work that utilizes both static and dynamic information but does not incur the overhead of running code at inference time.

B. Encoding Dynamic Information

Encoding dynamic behavior in the model has been studied extensively [45]–[48]. The most similar work to ours that we know of is Trex [49], which identifies semantically similar functions by encoding micro-execution traces during training and then transfers learning those values onto a pre-trained model. Despite the similar training/inference process there are several important differences between REINFORCEST and Trex: 1) Trex learns from binary level instruction sets, 2) during training it only maximizes agreement between similar sets and does not maximize distance between dissimilar samples, and 3) it would not work with proprietary models where transfer learning is not possible. GraphCodeBERT [50] is another prior work that we compare to directly with the experiments in Section V. Execution specific dataflow is a dynamic feature, but GraphCodeBERT extracts dataflow from the control flow graph without executing the code. Other similar works include Boost [51], a contrastive learning approach that generates semantically identical code and trains on the generated examples. This work identifies behaviorally similar code but does so without considering runtime information.

Learning from dynamic behavior is especially prevalent in the security community as malware often obfuscates source code to avoid detection. Most learning approaches that leverage runtime information function like EnDroid [52], a malware detection agent for android that learns from dynamic features

like execution traces and system activity, but requires collecting these dynamic features at inference time. Approaches like this take a considerably longer time to perform inference with since each sample must be examined relative to the query and runtime information must be collected for all queries. Another system, ScaleMalNet [53], mitigates these drawbacks by using distributed system techniques in training and pre-processing to scale their system, but still requires running the code at inference time.

VII. THREATS TO VALIDITY AND LIMITATIONS

A. Internal

We implemented benchmark techniques, including COSAL, using our unique experimental setting of training, validation, and testing splits. While recreating NLP-based search and AST diff was straightforward, COSAL’s snippet execution technique was more complex. Despite potential errors, our results are consistent with COSAL’s findings. We evaluated benchmarks using a stratified sample of 200 test queries, but performance might differ on the full dataset.

Using the Atcoder dataset, we aimed to compare with COSAL’s performance. While we executed as many samples as possible to compile the SSS, the lack of input specifications limited execution. We used default values when SSS couldn’t be collected, which might have affected the impact of SSS. Additionally, some pairs, labeled as true positives, had different input types (e.g., set vs. list), scoring an SSS of 0 during training. Finding a true SSS for these pairs could change the results.

We found that setting α to 0.2 had the best results of the values we tried, but more rigorous hyper-parameter tuning could result in a different optimal value.

B. External

The Atcoder dataset is compiled from submissions to a coding contest by practitioners of all levels. It remains unclear if the code in these contests is actually reflective of production code. Further we examine code-to-code search when two functions are defined to solve the exact same problem. It is unlikely to find such examples in practice and it remains unclear how well REINFORCEST would perform in real-world scenarios. As was mentioned earlier in Section IV we only consider behavioral clone, or not, during training and this limits the amount of context the model can learn from multiple positive and negative samples. Our experiments show the presence of both positive and negative samples impact performance but with a different training procedure that considered the distance between samples rather than classifying them as one of two categories the addition of more available samples during training may have more impact on performance.

C. Construct

We chose our evaluation metrics based on what we believed to be standard and meaningful. There are other important evaluation metrics like recall, overall accuracy, and F1 score that we did not measure. We chose cross-language search

because we believe that it was the most meaningful problem in identifying behavioral similarity, but experiments on same-language search would need to be carried out explicitly to evaluate REINFORCEST’s effectiveness for that purpose.

D. Limitations

Finding appropriate inputs that can be used to calculate the SSS between two pairs of functions is analogous to input (test case) generation which remains an open research problem. Further our method depends on a single function to function mapping which may not occur in practice, as the same behavior can be divided between multiple functions in different codebases.

VIII. CONCLUSION

This paper presents a novel code-to-code search technique and an implementation of that technique called REINFORCEST. The technique leverages modern LLM’s and enhances them by encoding runtime behavior in the form of a Semantic Similarity Score (SSS). Unlike other code search techniques that only consider positive samples during training, REINFORCEST both minimizes the distance between similar samples and maximizes the distance between dissimilar samples. We evaluated REINFORCEST on cross-language code search, both searching a corpus of Python samples with Java queries and vice versa. We found that our technique outperforms the state of the art on all evaluation metrics. Further, we found that our training technique always improves performance over modern LLM encodings out of the box, regardless of the underlying model, and we found that REINFORCEST performs especially well when the LLM models can be fine-tuned. This proves that our procedures presented here are the source of the improved performance and not the underlying models. We also present experiments quantifying both the impact of SSS and the importance of both positive and negative samples available during training. We found that including the SSS always improves performance, indicating the model does in fact learn dynamic behavior without dynamic context available at inference time, and that performance depends on the presence of both positive and negative samples. The complete set of results including results that were omitted for clarity is included in the supplementary material and all models and training procedures will be open-sourced.

REFERENCES

[1] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, J. Burstein, C. Doran, and T. Solorio, Eds. Association for Computational Linguistics, 2019, pp. 4171–4186. [Online]. Available: <https://doi.org/10.18653/v1/n19-1423>

[2] B. Vasilescu, A. Serebrenik, and V. Filkov, "The state of the art in distributed software development: A systematic review," in *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 2019, pp. 69–79.

[3] D. Kim, J. Nam, and J.-G. Song, "Mining libraries and frameworks for api migration," in *2017 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 2017, pp. 358–367.

[4] M. Zhou, Z. Zhang, J. Wu, H. Zhang, and H. Mei, "Automatic cross-language api adaptation via enhanced search-based api code transplantation," in *Proceedings of the 40th International Conference on Software Engineering: Companion Proceedings*, 2018, pp. 81–82.

[5] F. Long and M. Rinard, "Cross-language program repair using conditioned symbolic execution," in *2018 IEEE/ACM 40th International Conference on Software Engineering (ICSE)*. IEEE, 2018, pp. 738–748.

[6] D. Kim, J. Nam, and J.-G. Song, "Automatic patch generation learned from human-written patches," in *Proceedings of the 35th International Conference on Software Engineering*. IEEE/ACM, 2013, pp. 802–811.

[7] S. Yao, J. Huang, J. Chen, T. Xie, and X. Xia, "Exploring cross-language code search for programming education: A feasibility study," in *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*, 2019, pp. 259–265.

[8] R. Yue, Z. Gao, N. Meng, Y. Xiong, X. Wang, and J. D. Morgenthaler, "Automatic clone recommendation for refactoring based on the present and the past," in *2018 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, 2018, pp. 115–126.

[9] T. Yang and R. Fonseca, "Practical runtime specialization for higher-order functions," in *2021 USENIX Annual Technical Conference (USENIX ATC 21)*, 2021, pp. 601–614.

[10] G. Mathew and K. T. Stolee, "Cross-language code search using static and dynamic analyses," ser. ESEC/FSE 2021. New York, NY, USA: Association for Computing Machinery, 2021, p. 205–217. [Online]. Available: <https://doi.org/10.1145/3468264.3468538>

[11] K. Nakamura, R. Iwasaki, S. Hoshino, and I. Sato, "Atcoder: A dataset for machine learning on programming," in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 2021, pp. 3601–3611.

[12] Z. Feng, D. Guo, D. Tang, N. Duan, X. Feng, M. Gong, L. Shou, B. Qin, T. Liu, D. Jiang, and M. Zhou, "CodeBERT: A pre-trained model for programming and natural languages," in *Findings of the Association for Computational Linguistics: EMNLP 2020*, Nov. 2020, pp. 1536–1547.

[13] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, "Language models are few-shot learners," 2021.

[14] B. Ray, D. Kim, B. Vasilescu, S. Godhane, and V. Filkov, "Onward! toward more realistic program benchmarks." ACM, 2016, pp. 88–103.

[15] F.-H. Su, J. Bell, G. Kaiser, and S. Sethumadhavan, "Identifying Functionally Similar Code in Complex Codebases," in *2016 IEEE 24th International Conference on Program Comprehension (ICPC)*, 2016, pp. 1–10.

[16] Y. Kamei, A. Monden, and K.-i. Matsumoto, "Cross-language program search using a multilingual abstract syntax tree," in *Proceedings of the 2013 IEEE International Conference on Software Maintenance*. IEEE, 2013, pp. 320–329.

[17] L. Wu, T. Chen, and Y. Xie, "Cross-language code search with a bilingual semantic space," in *Proceedings of the 38th International Conference on Software Engineering Companion*. ACM, 2016, pp. 13–22.

[18] R. Xia, X. Xia, and J. Sun, "A neural model for cross-language code retrieval," in *Proceedings of the 40th International Conference on Software Engineering*. ACM, 2018, pp. 245–256.

[19] F.-H. Su, J. Bell, K. Harvey, S. Sethumadhavan, G. Kaiser, and T. Jebara, "Code Relatives: Detecting Similarly Behaving Software," ser. FSE 2016. New York, NY, USA: Association for Computing Machinery, 2016, p. 702–714. [Online]. Available: <https://doi.org/10.1145/2950290.2950321>

[20] A. Rastogi, R. S. Gu, K. Kumar, and S. Chandra, "Prose: Programming language semantics for large-scale code search," in *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*. ACM, 2015, pp. 1745–1757.

[21] C. S. Marcum, J. J. Haas, and J. Clause, "A comparative study of static and dynamic code analysis for bug detection," in *2010 IEEE International Conference on Software Maintenance*. IEEE, 2010, pp. 1–10.

[22] OpenAI, "Openai's text embeddings," 2023. [Online]. Available: <https://platform.openai.com/docs/guides/embeddings>

[23] V. Nitin, A. Saieva, B. Ray, and G. Kaiser, "Direct: A transformer-based model for decompiled variable name recovery," in *NLP4Prog 2021*, p. 48, 2021.

[24] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," *arXiv preprint arXiv:2002.05709*, 2020.

[25] G. Mathew, C. Parnin, and K. T. Stolee, "Slacc: Simion-based language agnostic code clones," in *2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE)*, 2020, pp. 210–221.

[26] J. Johnson, M. Douze, and H. Jégou, "Billion-scale similarity search with gpus," *IEEE Transactions on Big Data*, vol. 7, no. 3, pp. 535–547, 2019.

[27] C. M. Bishop, "Pattern recognition and machine learning," *Springer*, 2006.

[28] A. Moffat and J. Zobel, "Average rank performance metrics," *Journal of Information Retrieval*, vol. 11, no. 4, pp. 331–353, 2008.

[29] P. Willett, "Some issues in the evaluation of text retrieval systems," in *Proceedings of the seventeenth annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 1994, pp. 188–195.

[30] S. MacAvaney, "Bm25: A python implementation of the best match 25 information retrieval algorithm." [Online]. Available: https://github.com/dorianbrown/frank_bm25

[31] G. Mathew, C. Parnin, and K. T. Stolee, "Slacc: Simion-based language agnostic code clones," in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, 2020, pp. 210–221.

[32] K. Zhang, "A simple algorithm for tree edit distance," in *Proceedings of the 1st annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 1989, pp. 90–97.

[33] D. Guo, S. Lu, N. Duan, Y. Wang, M. Zhou, and J. Yin, "Unicoder: Unified cross-modal pre-training for code representation," *CoRR*, vol. abs/2203.03850, 2022. [Online]. Available: <https://doi.org/10.48550/arXiv.2203.03850>

[34] S. Lu, D. Guo, S. Ren, J. Huang, A. Svyatkovskiy, A. Blanco, C. Clement, D. Drain, D. Jiang, D. Tang *et al.*, "Codexglue: A machine learning benchmark dataset for code understanding and generation," *arXiv preprint arXiv:2102.04664*, 2021. [Online]. Available: <https://arxiv.org/abs/2102.04664>

[35] Microsoft, "Github," (Accessed on May 01 2020). [Online]. Available: www.github.com

[36] L. Di Grazia and M. Pradel, "Code search: A survey of techniques for finding code," *ACM Comput. Surv.*, vol. 55, no. 11, feb 2023. [Online]. Available: <https://doi.org/10.1145/3565971>

[37] T. Kamiya, S. Kusumoto, and K. Inoue, "Ccfinder: a multilingual token-based code clone detection system for large scale source code," *IEEE Transactions on Software Engineering*, vol. 28, no. 7, pp. 654–670, 2002.

[38] Z. Li, S. Lu, S. Myagmar, and Y. Zhou, "Cp-miner: finding copy-paste and related bugs in large-scale software code," *IEEE Transactions on Software Engineering*, vol. 32, no. 3, pp. 176–192, 2006.

[39] I. D. Baxter, A. Yahin, L. Moura, M. Sant'Anna, and L. Bier, "Clone detection using abstract syntax trees," in *Proceedings. International Conference on Software Maintenance (Cat. No. 98CB36272)*. IEEE, 1998, pp. 368–377.

[40] K. Kim, D. Kim, T. F. Bissyandé, E. Choi, L. Li, J. Klein, and Y. L. Traon, "Facoy: A code-to-code search engine," ser. ICSE '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 946–957. [Online]. Available: <https://doi.org/10.1145/3180155.3180187>

[41] S. N. Pinku, D. Mondal, and C. K. Roy, “Pathways to leverage transcompiler based data augmentation for cross-language clone detection,” 2023.

[42] Y. Zou, B. Ban, Y. Xue, and Y. Xu, “Cggraph: A pdg-based code clone detector with approximate graph matching,” ser. ASE ’20. New York, NY, USA: Association for Computing Machinery, 2021, p. 931–942. [Online]. Available: <https://doi.org/10.1145/3324884.3416541>

[43] H. Sajnani, V. Saini, J. Svajlenko, C. K. Roy, and C. V. Lopes, “Sourcererc: Scaling code clone detection to big-code,” ser. ICSE ’16. New York, NY, USA: Association for Computing Machinery, 2016, p. 1157–1168. [Online]. Available: <https://doi.org/10.1145/2884781.2884877>

[44] F.-H. Su, J. Bell, G. Kaiser, and B. Ray, “Obfuscation Resilient Search through Executable Classification,” ser. MAPL 2018. New York, NY, USA: Association for Computing Machinery, 2018, p. 20–30. [Online]. Available: <https://doi.org/10.1145/3211346.3211352>

[45] B. He and J. Liu, “Deeplog: Anomaly detection and diagnosis from system logs through deep learning,” in *2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2016, pp. 230–241.

[46] M. Allamanis, M. Brockschmidt, and M. Khademi, “Learning to represent programs with graphs,” in *International Conference on Learning Representations*, 2018.

[47] M. Watter, J. Springenberg, J. Boedecker, and M. Riedmiller, “Action-conditional video prediction using deep networks in atari games,” in *Advances in Neural Information Processing Systems*, 2015, pp. 2863–2871.

[48] H. Pham, T. Tran, and S. Venkatesh, “Modelling dynamic textual content with temporal language models,” *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 6, no. 3, p. 33, 2015.

[49] K. Pei, Z. Xuan, J. Yang, S. Jana, and B. Ray, “Learning Approximate Execution Semantics From Traces for Binary Function Similarity,” *IEEE Transactions on Software Engineering*, vol. 49, no. 4, pp. 2776–2790, April 2023. [Online]. Available: <https://doi.org/10.1109/TSE.2022.3231621>

[50] D. Guo, S. Ren, S. Lu, Z. Feng, D. Tang, S. Liu, L. Zhou, N. Duan, J. Yin, D. Jiang *et al.*, “Graphcodebert: Pre-training code representations with data flow,” in *International Conference on Learning Representations*, 2021.

[51] Y. Ding, L. Buratti, S. Chakraborty, S. Pujar, A. Morari, and B. Ray, “Contrastive learning for source code with structural and functional properties,” 2022. [Online]. Available: <https://openreview.net/forum?id=7KgeqhkbZab>

[52] P. Feng, J. Ma, C. Sun, X. Xu, and Y. Ma, “A novel dynamic android malware detection system with ensemble learning,” *IEEE Access*, vol. 6, pp. 30 996–31 011, 2018.

[53] R. Vinayakumar, M. Alazab, K. P. Soman, P. Poornachandran, and S. Venkatraman, “Robust intelligent malware detection using deep learning,” *IEEE Access*, vol. 7, pp. 46 717–46 738, 2019.