

Towards Better Code Understanding in Decoder-Only Models with Contrastive Learning

Jiayi Lin^{1,2}, Yanlin Wang^{3*}, Yibiao Yang⁴, Lei Zhang¹, Yutao Xie¹

¹International Digital Economy Academy, Shenzhen, China

²Tencent, Shenzhen, China

³School of Software Engineering, Sun Yat-sen University, Zhuhai, China

⁴State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing, China

jiayilin1024@gmail.com, yanlin-wang@outlook.com, yangyibiao@nju.edu.cn, {leizhang,yutaoxie}@idea.edu.cn

Abstract

Recent advances in large-scale code generation models have led to remarkable progress in producing high-quality code. These models are trained in a self-supervised manner on extensive unlabeled code corpora using a decoder-only architecture. However, despite their generative strength, decoder-only models often exhibit limited performance on code understanding tasks such as code search and clone detection, primarily due to their generation-oriented training objectives. While training large encoder-only models from scratch on massive code datasets can improve understanding ability but remains computationally expensive and time-consuming. In this paper, we explore a more efficient alternative by transferring knowledge from pre-trained decoder-only code generation models to code understanding tasks. We investigate how decoder-only architectures can be effectively adapted to learn discriminative and semantically meaningful code representations. To this end, we propose CL4D, a contrastive learning framework tailored to strengthen the representation capabilities of decoder-only models. Extensive experiments on multiple benchmark datasets demonstrate that CL4D achieves competitive or superior performance compared to existing methods on representative code understanding tasks, including code search and clone detection. Further analysis reveals that CL4D substantially improves the semantic alignment of code representations by reducing the distance between semantically similar code snippets. These findings highlight the feasibility of leveraging decoder-only models as a unified backbone for both code generation and understanding.

Code — <https://github.com/JiayiLin1024/CL4D>

Introduction

Recent years have witnessed a surge of large-scale pre-trained code generation models. These models are trained in a self-supervised manner on vast code corpora using a decoder-only architecture and have demonstrated remarkable success (Nijkamp et al. 2023; Guo et al. 2024; Lozhkov et al. 2024). They have been widely adopted in developer-assistance tools, such as GitHub Copilot¹ and Cursor², to support everyday programming tasks.

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://github.com/features/copilot>

²<https://cursor.com/>




Figure 1: Timeline of code-large models. For visual clarity, only the largest size of representative models is plotted. The orange arrow indicate the development trends of encoder-only code understanding models, while the green arrow represent the trends of decoder-only code generation models.

Despite their outstanding generative performance, decoder-only models often struggle with code understanding tasks, including code search and code clone detection. This limitation arises from their auto-regressive training objectives, which primarily emphasize next-token prediction rather than semantic comprehension. Previous studies have shown that stronger performance on code understanding tasks can be achieved by pre-training encoder-only models from scratch on large-scale code datasets (Hoffmann et al. 2022; Lin et al. 2024). However, such training is computationally expensive and time-consuming, making it impractical for widespread adoption. As illustrated in Figure 1, the landscape of code models has evolved rapidly: from 1.1B SantaCoder (Allal et al. 2023) and the 15B StarCoder (Li et al. 2023) to the 70B CodeLlama (Rozière et al. 2023). These models are significantly larger and more data-rich than the current leading code understanding models, such as the 1.3B CodeSage (Zhang et al. 2024) and the 1.5B CoLSBERT (Lin et al. 2024). Nevertheless, how to effectively leverage these powerful pre-trained decoder-only code generation models for code understanding tasks remains an open and unexplored problem.

In this paper, we investigate the adaption of pre-trained

decoder-only code generation models to code understanding tasks. We explore two methods for extracting code representations from the decoder-only Transformer: 1) using the encoding of the last token, and 2) using the average encodings of all tokens. With these code representations, we introduce CL4D, a contrastive learning framework that continues pre-training the code generation models to enhance their representational capabilities. Follow Gao et al. (Gao et al. 2021), we employ hard negatives, semantically distinct code snippets that are close to the query in representation space, to improve training efficiency.

To ensure rigorous evaluation, we compare pre-trained decoder-only and encoder-only models of comparable size on two representative understanding tasks: code search and clone detection. Experimental results indicate that our proposed CL4D effectively adapts decoder-only code generation models for code understanding tasks. After continued pre-training on a small code corpus, these models outperform pre-trained encoder-only models on most tasks. Furthermore, we observe that larger decoder-only models further improve code understanding performance.

In summary, the contributions of this paper are as follows:

- We investigate how pre-trained decoder-only code generation models can be effectively leveraged for code understanding tasks, bridging the gap between generation-oriented and understanding-oriented applications.
- We propose CL4D, a contrastive learning–based continued pre-training approach that enhances the representation ability of decoder-only models, enabling competitive or superior performance to encoder-only models on code understanding benchmarks.
- Our findings reveal that larger decoder-only models benefit code understanding tasks and demonstrate the potential for unifying code understanding and code generation within a single model architecture.

Background and Motivation

In recent years, large-scale language models with a decoder-only architecture have achieved remarkable success in the field of code generation. Models such as Codex (Chen et al. 2021), CodeGen (Nijkamp et al. 2023), StarCoder (Li et al. 2023), and Code Llama (Rozière et al. 2023) have demonstrated impressive capabilities in generating syntactically correct and semantically meaningful code. These models are typically pre-trained on massive corpora of natural language and programming languages, allowing them to acquire extensive programming knowledge and multi-language fluency. With the emergence of instruction-tuned variants like WaveCoder (Yu et al. 2024) and CodeGemma (Zhao et al. 2024), decoder-only models continue to push the boundaries of generative performance.

Despite their superiority in generative tasks, our empirical observations reveal a surprising performance gap when these models are applied directly to code understanding tasks such as code search and code clone detection. As shown in Table 4, encoder-only models—such as GraphCodeBERT (Guo et al. 2021), UniXcoder (Guo et al. 2022)

and CodeSage (Zhang et al. 2024)—which are explicitly designed for understanding source code, consistently outperform their decoder-only counterparts on benchmark code understanding datasets.

This performance gap persists even though decoder-only models possess several natural advantages. First, they have been exposed to significantly larger and more diverse code corpora during pre-training, often spanning multiple programming languages and domains. Second, these models are frequently trained with objectives aligned with next-token prediction, which implicitly teaches structural and semantic patterns in code. However, their lack of bi-directional attention mechanisms and their generative training objective may limit their capacity for fine-grained code comprehension.

One obvious solution is to pre-train a new encoder-only model at the same scale as the latest decoder-only models. However, the computational and financial costs of such an endeavor are prohibitively high, particularly when targeting billion-parameter models trained over trillions of tokens. This raises a critical research question: Can we enhance the code comprehension capabilities of existing decoder-only models without incurring the massive cost of training new models from scratch?

To this end, we propose a contrastive learning framework aimed at improving the code understanding performance of decoder-only models. Our approach leverages the inherent strengths of pre-trained decoder-only models while avoiding the prohibitive cost of full model retraining.

Methodology

In this section, we explore obtaining encoding from decoder-only pre-trained code generation models to address code understanding tasks. The overview of CL4D is shown in Figure 2. We first introduce the data construction method and then introduce the training of CL4D, including the model architecture and contrastive learning.

Data Construction

We train decoder-only Transformer models on six programming languages from The Stack dataset (Kocetkov et al. 2022): Python, Java, Go, PHP, JavaScript, and Ruby. To enhance the ability of model to differentiate between natural language and programming language, we use Tree-Sitter³ to extract bimodal data from The Stack dataset for training, i.e., (*query*, *code*) pairs. The query is the first line of the function docstring, which approximates the query of user. Additionally, we apply the filtering rules to improve the quality of the training corpus following CSN (Husain et al. 2019). Given that pre-trained code generation models have already learned a vast amount of code knowledge, we construct millions of data samples for continued pre-training.

Model Architecture

We investigated two methods for obtaining code representation from decoder-only models: 1) using the embedding of the last token, and 2) averaging the embeddings of all

³<https://tree-sitter.github.io/tree-sitter/>

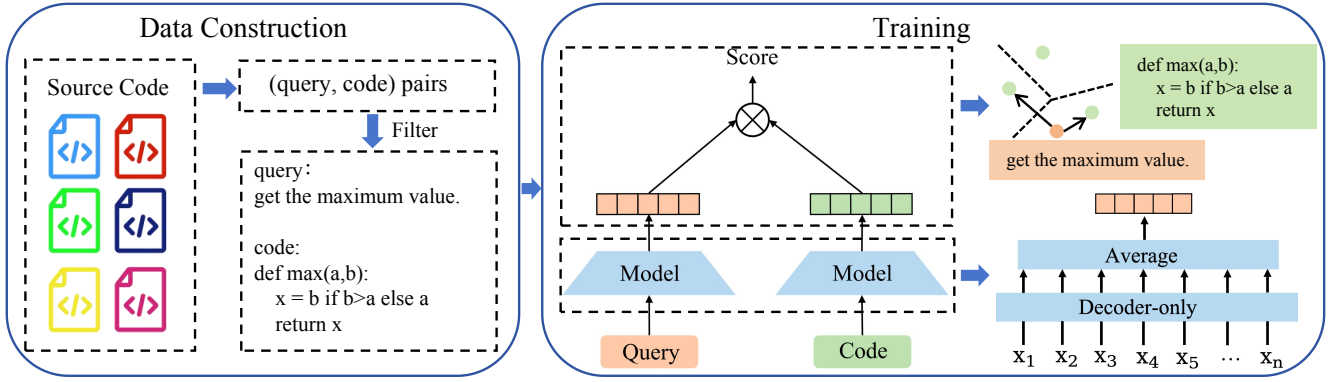


Figure 2: Overview of CL4D.

tokens. Encoder-only models like BERT typically use the embedding of the first token $[CLS]$ in the last hidden layer as the feature representation of the input code snippet, utilizing a bidirectional attention mechanism. By contrast, decoder-only models use a unidirectional attention mechanism, where only subsequent tokens can refer to previous information. Consequently, only the last token aggregates the information of the entire sample. In the first method, we use the embedding of the last token in the last hidden layer as the feature representation of the input code snippet. In the second method, we calculate the average of the embeddings of all input tokens in the last hidden layer as the feature representation.

Following previous work (Guo et al. 2021; Zhang et al. 2024), we adopt a dual-encoder architecture to continue pre-training the code encoding model. As depicted in Figure 2, this architecture consists of two Transformer decoder modules with shared weights to encode the input code snippet. In our experiments, we initialize the Transformer decoder module with the pre-trained code generation models.

Contrastive Learning

The traditional left-to-right training method for predicting the next token limits the representation capability of code generation models. To fully leverage the extensive code knowledge in the existing open-source code generation models and to avoid the high costs of retraining code understanding models, we employ contrastive learning to further fine-tune these models, thereby improving their representation abilities.

Contrastive learning is the primary method for enhancing the representation ability of models (Zhang, Lan, and He 2023). Here, we employ contrastive learning tasks to train pre-trained code generation models, addressing the limited representation ability of decoder-only Transformer due to the unidirectional attention mechanism.

Specifically, we randomly sample n samples from the training corpus to form a batch of data $B = (q_1, c_1), \dots, (q_n, c_n)$. For a (q_i, c_i) pair, we take c_i as the positive sample c^+ for q_i , and the other $n - 1$ codes in the batch, excluding c_i , as negative samples for q_i . Additionally,

following Gao et al. (Gao et al. 2021), we construct hard negative samples c^h for q_i . Specifically, during the data processing stage, we utilize UniXcoder to select a code snippet from the entire codebase for each query in the training set, i.e., the code snippet that is close to the query in the representation space but have different semantic from the query. The loss function can be formulated as:

$$\mathcal{L} = -\log \frac{\exp(s(q, c^+)/\tau)}{\sum_{i=1}^n \exp(s(q, c_i)/\tau) + \exp(s(q, c^h)/\tau)}, \quad (1)$$

where τ is the temperature hyper-parameter, which we set to 0.05. $s(q, c)$ denotes the correlation score between the two, calculated using cosine similarity. The training objective is to minimize the distance between related pairs of samples in the representation space while maximizing the distance between unrelated pairs. Each batch is randomly sampled from the entire corpus, allowing for mixed programming languages within a batch, which helps the pre-trained model learn a unified semantic space for codes in different programming languages.

Experimental Setup

In this section, we conduct experiments to address the following research questions:

- RQ1: Can our proposed CL4D boost the performance of the decoder-only code generation models for the code understanding tasks?
- RQ2: Which method is more suitable for a decoder-only model to extract code representation?
- RQ3: To what extent does contrastive learning enhance the code representations of decoder-only models?
- RQ4: How does CL4D perform in the zero-shot setting?
- RQ5: Why does our approach work?

Understanding Tasks

We evaluate the code understanding capabilities of decoder-only Transformer on two primary tasks:

- **Code Search:** This task aims to find the most semantically similar code snippet in a candidate codebase based

on a given natural language query. We perform experiments on two datasets: CodeSearchNet (CSN) (Husain et al. 2019) and CoSQA (Huang et al. 2021). CSN is a multi-language code corpus collected from GitHub, including Python, Java, Go, PHP, JavaScript, and Ruby. CoSQA contains real user queries from Microsoft Bing search logs, with code snippets from CSN. We use Mean Reciprocal Rank (MRR) as the evaluation metric, which is widely adopted in previous studies (Guo et al. 2022; Zhang et al. 2024).

- **Clone Detection:** This task aims to identify the most semantically similar code snippet in a candidate code-base based on a given code snippet as query. We conduct experiments on the POJ-104 dataset (Mou et al. 2016), designed to measure semantic similarity between code snippets. Mean Average Precision (MAP) is the evaluation metric for this task.

Baselines

We compare the code understanding capabilities of encoder-only models and decoder-only models of comparable sizes, and evaluate the extent to which the proposed CL4D method enhances the performance of decoder-only models.

For encoder-only models, we select four state-of-the-art pre-trained models:

- **CodeBERT** (Feng et al. 2020) (125M) is trained with masked language modeling and replaced token detection on six programming languages in CodeSearchNet.
- **GraphCodeBERT** (Guo et al. 2021) (125M) is trained with additional code structure understanding tasks, including edge prediction and node alignment.
- **UniXcoder** (Guo et al. 2022) (125M) is trained with multiple language modeling and contrastive learning tasks. It is an encoder-decoder architecture model, where only the encoder is used for code understanding tasks via a special token `< encoder_only >`.
- **CodeSage** (Zhang et al. 2024) (1.3B) is trained with masked language modeling and contrastive learning on the extensive code corpus from The Stack dataset.

For decoder-only models, we select four state-of-the-art pre-trained models:

- **CodeGPT** (Lu et al. 2021) (125M) is trained on CodeSearchNet with the same structure and training objectives as GPT-2 (Radford et al. 2019).
- **CodeGen** (Nijkamp et al. 2023) (350M) is a model for program analysis trained on extensive natural language and programming language data. We choose the 350M version to facilitate comparison with similarly sized encoder-only models.
- **SantaCoder** (Allal et al. 2023) (1.1B) is trained with multi-query attention and fill-in-the-middle on Java, JavaScript, and Python subsets of The Stack, incorporating near-deduplication and comments-to-code filtering.
- **phi-1** (Gunasekar et al. 2023) (1.3B) is a code generation model trained using textbook data sourced from the internet and textbooks synthesized by GPT-3.5.

- **DeepSeek-Coder** (Guo et al. 2024) (1.3B) is a model trained on a high-quality project-level code corpus.

Training details

During the pre-training phase, we utilize the tokenizer from each decoder-only model directly for tokenization. These models are trained using the AdamW optimizer, with a learning rate set to $2e-5$. We train the decoder-only models for 2 epochs on the dataset, employing a batch size of 64. The experiments are conducted on 8 GPUs (80G A100). The longest pre-training duration is observed for phi-1, which takes approximately 3 days. In the fine-tuning phase, we adhere to the experimental settings of UniXcoder to minimize the impact of hyper-parameters on the fine-tuning results.

Evaluation

RQ1: Overall performance comparison

We evaluate the effectiveness of CL4D on two representative code understanding tasks: code search and clone detection. The experiment results are shown in Table 1, we can find that CL4D enables the decoder-only model to achieve state-of-the-art performance. Specifically, our method improves the performance of decoder-only models by approximately 2% compared to encoder-only models of the same size. Additionally, as the size of the decoder-only model increases, its performance on downstream code understanding tasks shows a corresponding improvement. Overall, our approach effectively enhances the code understanding capabilities of decoder-only models.

RQ2: Code representation extraction method study for decoder-only models

To identify the most effective strategy for extracting code representations from decoder-only Transformer models, we investigated two approaches: decoder-only last token (**Last**) and decoder-only average (**Avg**). In addition, we examined how different padding strategies influence the quality and effectiveness of the resulting code representations.

Using CodeGen (350M) as an example, we conducted an ablation experiment on the downstream task of code search, with results presented in Table 2. Our findings indicate that the optimal method for obtaining code representations varies with different padding strategies. When right padding is used, averaging all token embeddings yields better results than using only the embedding of the last token. Conversely, with left padding, the embedding of the last token outperforms the average of all embeddings. Considering results from both the CSN and CoSQA datasets, the decoder-only Transformer achieves the best representation with right padding and averaging all token embeddings. This setting was used in all other experiments.

RQ3: Ablation study

To investigate the effectiveness of our method, we designed ablation experiments to analyze the contributions of different components in CL4D.

We trained three groups of models based on the CodeGen model: 1) the original Codegen; 2) the model trained using

	Method	Code Search		Clone Detection
		CSN	CoSQA	POJ-104
Encoder-only	CodeBERT (Feng et al. 2020) (125M)	70.18	65.7	83.79
	GraphCodeBERT (Guo et al. 2021) (125M)	72.08	68.4	85.50
	UniXcoder (Guo et al. 2022) (125M)	74.40	70.1	89.56
	CodeSage (Zhang et al. 2024) (1.3B)	75.80	68.0	87.70
Decoder-only	CodeGPT (Lu et al. 2021) (125M) + CL4D	70.20	69.0	87.96
	CodeGen (Nijkamp et al. 2023) (350M) + CL4D	73.30	71.5	89.68
	SantaCoder (Allal et al. 2023) (1.1B) + CL4D	74.98	72.2	83.98
	phi-1 (Gunasekar et al. 2023) (1.3B) + CL4D	75.18	72.8	92.72
	DeepSeek-Coder (Guo et al. 2024) (1.3B) + CL4D	77.57	71.9	89.71

Table 1: Comparison of the performance of encoder-only models and decoder-only models on code understanding tasks.

Type	CSN							CoSQA
	Python	Java	Go	Php	JavaScript	Ruby	Avg	
Left padding + Avg	68.0	67.7	89.1	62.9	59.5	62.2	68.23	66.0
Right padding + Avg	69.0	70.6	90.1	64.4	64.8	69.6	71.42	71.6
Left padding + Last	70.1	71.0	90.1	64.9	65.7	70.1	71.98	67.9
Right padding + Last	67.4	70.1	90.2	54.7	64.1	56.1	67.10	68.9

Table 2: The performance of code representations acquired in different ways on code search tasks. Bold indicates optimal performance, and bold and underlined indicates suboptimal performance.

the contrastive learning method with in-batch negatives; 3) the model trained using the contrastive learning method with in-batch negatives and hard negatives. We compared their zero-shot performance on downstream code understanding tasks, with results shown in Table 3. The findings reveal that using in-batch negatives significantly enhances the code understanding ability of decoder-only models. Adding hard negatives further improves performance by approximately 1.5%. Contrastive learning is expected to become an effective method for the decoder-only Transformer architecture to unify code understanding and code generation tasks.

RQ4: Zero-shot

We evaluated baseline models without fine-tuning on code search and clone detection tasks. As shown in Table 4, the zero-shot performance of encoder-only models on two downstream code understanding tasks improves with the introduction of more code structure-related tasks or scaling up strategies. Notably, UniXcoder and CodeSage leverage contrastive learning tasks during pre-training, significantly enhancing zero-shot performance compared to other encoder-only models. By contrast, decoder-only models underperform on downstream tasks related to code understanding, as their pre-training primarily focuses on next-token prediction rather than bidirectional semantic representation.

The experimental results indicate that the extensive

Method	Code Search		Clone Detection
	CSN	CoSQA	POJ-104
CL4D	72.00	51.20	45.84
- Hard Negative	70.80	50.40	44.65
- In-Batch Negative	1.42	0.45	13.20

Table 3: Ablation results of CL4D in the zero-shot setting.

code knowledge acquired by pre-trained code understanding models remains underutilized. To leverage this knowledge for code understanding tasks, we applied the CL4D method described in Section for continued pre-training of these decoder-only models. The experimental outcomes, presented in Table 4, demonstrate that: 1) Our proposed method significantly enhances the zero-shot performance of decoder-only models on both code search and clone detection tasks, with average improvements ranging from 40% to 60%, and a maximum improvement of 75.90%. 2) Our method enables the zero-shot performance of decoder-only models on various code understanding tasks to surpass that of encoder-only models of comparable size. 3) Our method enables decoder-only models to match the performance of same-sized, fine-tuned encoder-only state-of-the-art models on the CSN dataset, even without fine-tuning.

RQ5: Why does our approach work?

Case study Figure 3 displays the top-1 ranked code snippets returned by DeepSeek-Coder, CodeSage, and DeepSeek-Coder+CL4D for the query “python only list files with specific extension”. DeepSeek-Coder+CL4D successfully returns the correct code snippet, which iterates through files and identifies those matching the specified extension. In contrast, the code snippets from DeepSeek-Coder and CodeSage are incorrect: DeepSeek-Coder returns a result entirely unrelated to the query, while CodeSage’s snippet only identifies the extension without returning the corresponding files. These results indicate that continued pre-training of decoder-only models with CL4D significantly enhances their code understanding ability, allowing them to outperform encoder-only models of the same size.

Visualization To investigate why CL4D effectively improves the performance of decoder-only models on code un-

Method	Code Search		Clone Detection	
	CSN	CoSQA	POJ-104	
Encoder-only	CodeBERT (Feng et al. 2020) (125M)	0.10	0.24	20.38
	GraphCodeBERT (Guo et al. 2021) (125M)	11.26	16.20	31.27
	UniXcoder (Guo et al. 2022) (125M)	46.40	42.11	42.08
	CodeSage (Zhang et al. 2024) (1.3B)	71.24	47.53	73.07
Decoder-only	CodeGPT (Lu et al. 2021) (125M)	0.12	0.04	9.41
	CodeGen (Nijkamp et al. 2023) (350M)	1.42	0.45	13.20
	SantaCoder (Allal et al. 2023) (1.1B)	0.07	0.11	15.57
	phi-1 (Gunasekar et al. 2023) (1.3B)	6.66	6.75	27.62
	DeepSeek-Coder (Guo et al. 2024) (1.3B)	0.12	0.63	16.51
Decoder-only	CodeGPT (Lu et al. 2021) (125M) + CL4D	<u>67.56</u> (\uparrow 67.44)	<u>53.49</u> (\uparrow 53.45)	<u>25.93</u> (\uparrow 16.52)
	CodeGen (Nijkamp et al. 2023) (350M) + CL4D	71.97 (\uparrow 70.55)	51.18 (\uparrow 50.73)	45.84 (\uparrow 32.64)
	SantaCoder (Allal et al. 2023) (1.1B) + CL4D	74.18 (\uparrow 74.11)	52.82 (\uparrow 52.71)	71.14 (\uparrow 55.57)
	phi-1 (Gunasekar et al. 2023) (1.3B) + CL4D	73.21 (\uparrow 66.55)	52.95 (\uparrow 46.20)	69.99 (\uparrow 42.37)
	DeepSeek-Coder (Guo et al. 2024) (1.3B) + CL4D	76.02 (\uparrow 75.90)	48.34 (\uparrow 47.71)	71.18 (\uparrow 54.67)

Table 4: Comparison of the performance of encoder-only models and decoder-only models on code understanding tasks in the zero-shot settings. Bold indicates optimal performance, and bold and underlined indicates suboptimal performance.

Query: python only list files with specific extension.		
DeepSeek-Coder: <pre>def test3(): """Test the multiprocessing""" import time p = MVisionProcess() p.start() time.sleep(5) p.stop()</pre>	CodeSage: <pre>def watched_extension(extension): """Return True if the given extension is one of the watched extensions""" for ext in hamipy.VALID_EXTENSIONS: if extension.endswith('.' + ext): return True return False</pre>	DeepSeek-Coder+CL4D: <pre>def glob_by_extensions(directory, extensions): """ Returns files matched by all extensions in the extensions list """ directorycheck(directory) files = [] xt = files.extend for ex in extensions: xt(glob.glob('{0}/*.{1}'.format(directory, ex))) return files</pre>

Figure 3: Case study on the code search task in the zero-shot setting.

derstanding tasks, we visualized the representation distribution of CodeGPT and CodeGPT+CL4D across the CSN, CoSQA, and POJ-104 datasets. For the CSN dataset, each query has only one correct answer, we randomly sampled 30 query-code pairs from the test set to form 30 categories, each containing 2 samples. The same approach was applied to the CoSQA dataset. For the POJ-104 test set, which has 24 problems, each with multiple correct answers, we sampled 5 correct answers per problem, forming 24 categories with 5 samples each.

We then feed the sampled data into the models to obtain their high-dimensional feature space representations. Using t-SNE (Van der Maaten and Hinton 2008), we reduce the dimensionality of these representations and visualize their distribution, as shown in Figure 4. Points with the same color and shape represent samples with the same semantics. Ideally, their representations in the vector space should be as close as possible. The results show that the representations encoded by the original decoder-only model are highly scattered, with semantically similar samples far apart in the vector space. The decoder-only model pre-trained with CL4D aligns the representations of semantically similar samples,

bringing them closer together in the semantic space.

Overall, the visualization clearly demonstrates that our method significantly enhances the representation capability of the decoder-only models.

Related works

Code Representation Recent work has focused on obtaining programming language representations for code understanding tasks. Models like CodeBERT (Feng et al. 2020), StarEncoder (Li et al. 2023), CodeSage (Zhang et al. 2024), and CoLSBERT (Lin et al. 2024) use the Masked Language Modeling (MLM) objective from NLP, applying it to serialized code. Recognizing the structural differences between code and text, some models enhance learning by incorporating code structures. For example, GraphCodeBERT (Guo et al. 2021) uses data flow, TreeBERT (Jiang et al. 2021) utilizes the abstract syntax tree (AST), and SynCoBERT (Wang et al. 2021) and UniXcoder (Guo et al. 2022) include serialized ASTs during training.

Although these models advance code understanding, they remain relatively small, typically trained on datasets with only millions of samples, and even the largest model, CoLS-

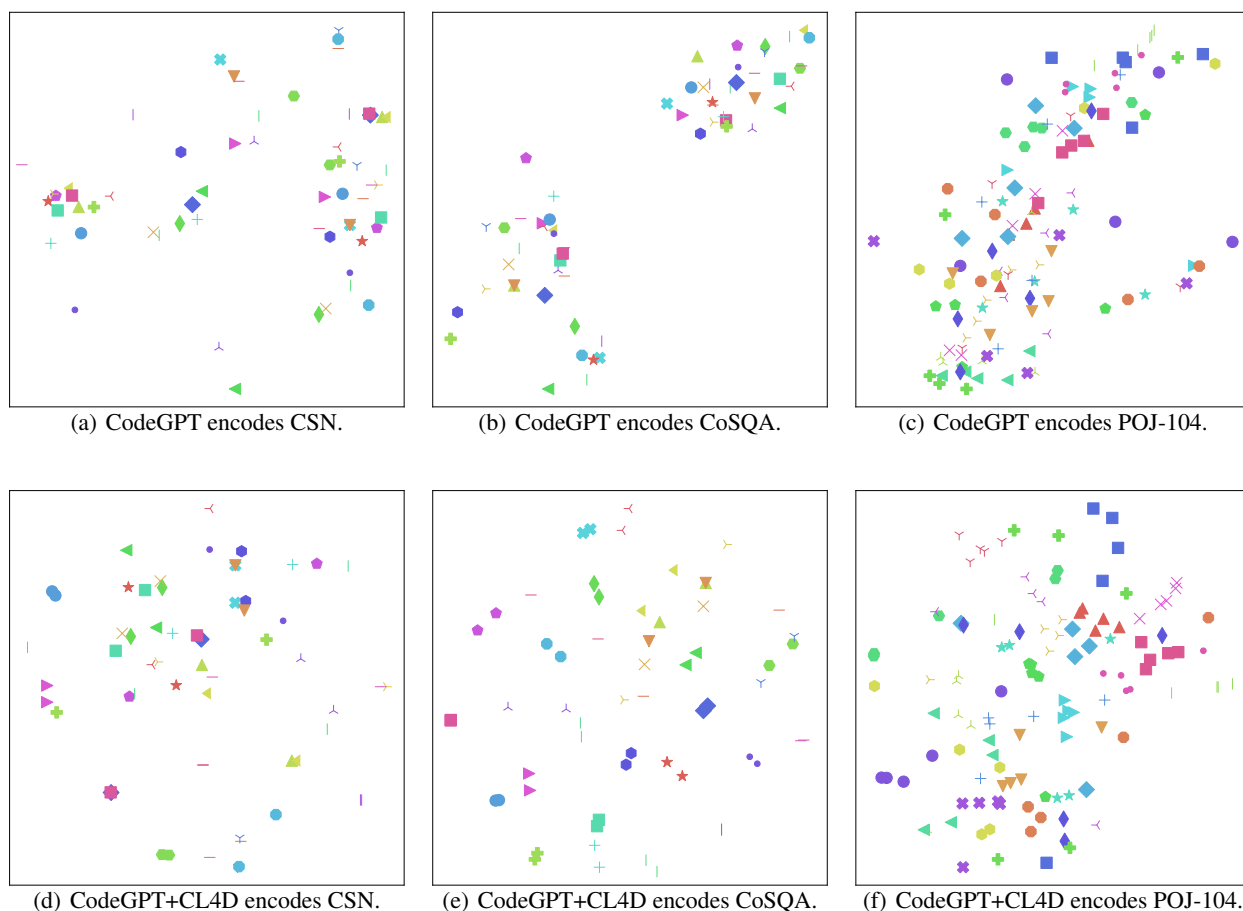


Figure 4: Visualization of code representations. The first row shows the encoding results of CodeGPT, while the second row illustrates the encoding results of CodeGPT+CL4D.

BERT, has only 1.5B parameters. While scaling up model size and training data can further improve performance, it also incurs substantial computational and resource costs. Inspired by the success of large-scale decoder-only code generation models, our work aims to leverage such pre-trained models to improve code understanding tasks, achieving state-of-the-art performance.

Contrastive Learning for Code Contrastive learning has been widely employed in representation learning for deep learning networks. Recent studies have extended this paradigm to code corpora, proposing various methods to construct positive or negative samples to enhance code representation. For example, CoSQA (Huang et al. 2021) constructs positive samples by rewriting the query, SynCoBERT (Wang et al. 2021) and Code-MVP (Wang et al. 2022) create positive sample pairs across different code modalities using diverse information obtained during program compilation, and UniXcoder (Guo et al. 2022), inspired by SimCLR (Chen et al. 2020), uses dropout to construct positive samples in the latent vector space. Additionally, ChatDANCE (Wang et al. 2023) uses large language models such as ChatGPT to synthesize positive code samples. CodeRetriever (Li et al. 2022), R2 (Dong et al. 2023),

and CodeSage (Zhang et al. 2024) construct hard negatives for queries using rankers. In this paper, to mitigate the limited representation capability of decoder-only models by their unidirectional attention mechanism, we select code snippets with different semantics from the query but close in the representation space as hard negatives.

Conclusion

In this paper, we have presented a novel application of pre-trained decoder-only code generation models to code understanding tasks. We explore four distinct methods by which these decoder-only models develop code representations. Our study includes a comparative analysis of encoder-only and decoder-only models, examining their performance on downstream code understanding tasks in a zero-shot setting. To enhance the representational capabilities of the decoder-only model, we introduce a contrastive learning method, CL4D, which constructs in-batch negatives and hard negatives for training. Experimental results demonstrate that CL4D enables decoder-only models to significantly outperform encoder-only models in most tasks. Our findings indicate that decoder-only Transformer have the potential to unify code understanding and generation tasks.

References

- Allal, L. B.; Li, R.; Kocetkov, D.; Mou, C.; Akiki, C.; Ferrandis, C. M.; Muennighoff, N.; Mishra, M.; Gu, A.; Dey, M.; Umapathi, L. K.; Anderson, C. J.; Zi, Y.; Lamy-Poirier, J.; Schoelkopf, H.; Troshin, S.; Abulkhanov, D.; Romero, M.; Lappert, M.; Toni, F. D.; del Río, B. G.; Liu, Q.; Bose, S.; Bhattacharyya, U.; Zhuo, T. Y.; Yu, I.; Villegas, P.; Zocca, M.; Mangrulkar, S.; Lansky, D.; Nguyen, H.; Contractor, D.; Villa, L.; Li, J.; Bahdanau, D.; Jernite, Y.; Hughes, S.; Fried, D.; Guha, A.; de Vries, H.; and von Werra, L. 2023. SantaCoder: don't reach for the stars! *CoRR*, abs/2301.03988.
- Chen, M.; Tworek, J.; Jun, H.; Yuan, Q.; de Oliveira Pinto, H. P.; Kaplan, J.; Edwards, H.; Burda, Y.; Joseph, N.; Brockman, G.; Ray, A.; Puri, R.; Krueger, G.; Petrov, M.; Khlaaf, H.; Sastry, G.; Mishkin, P.; Chan, B.; Gray, S.; Ryder, N.; Pavlov, M.; Power, A.; Kaiser, L.; Bavarian, M.; Winter, C.; Tillet, P.; Such, F. P.; Cummings, D.; Plappert, M.; Chantzis, F.; Barnes, E.; Herbert-Voss, A.; Guss, W. H.; Nichol, A.; Paino, A.; Tezak, N.; Tang, J.; Babuschkin, I.; Balaji, S.; Jain, S.; Saunders, W.; Hesse, C.; Carr, A. N.; Leike, J.; Achiam, J.; Misra, V.; Morikawa, E.; Radford, A.; Knight, M.; Brundage, M.; Murati, M.; Mayer, K.; Welinder, P.; McGrew, B.; Amodei, D.; McCandlish, S.; Sutskever, I.; and Zaremba, W. 2021. Evaluating Large Language Models Trained on Code. *CoRR*, abs/2107.03374.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. E. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, 1597–1607. PMLR.
- Dong, H.; Lin, J.; Leng, Y.; Chen, J.; and Xie, Y. 2023. Retriever and Ranker Framework with Probabilistic Hard Negative Sampling for Code Search. *CoRR*, abs/2305.04508.
- Feng, Z.; Guo, D.; Tang, D.; Duan, N.; Feng, X.; Gong, M.; Shou, L.; Qin, B.; Liu, T.; Jiang, D.; and Zhou, M. 2020. CodeBERT: A Pre-Trained Model for Programming and Natural Languages. In Cohn, T.; He, Y.; and Liu, Y., eds., *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, 1536–1547. Association for Computational Linguistics.
- Gao, L.; Zhang, Y.; Han, J.; and Callan, J. 2021. Scaling Deep Contrastive Learning Batch Size under Memory Limited Setup. In Rogers, A.; Calixto, I.; Vulic, I.; Saphra, N.; Kassner, N.; Camburu, O.; Bansal, T.; and Shwartz, V., eds., *Proceedings of the 6th Workshop on Representation Learning for NLP, RePLANLP@ACL-IJCNLP 2021, Online, August 6, 2021*, 316–321. Association for Computational Linguistics.
- Gunasekar, S.; Zhang, Y.; Aneja, J.; Mendes, C. C. T.; Giorno, A. D.; Gopi, S.; Javaheripi, M.; Kauffmann, P.; de Rosa, G.; Saarikivi, O.; Salim, A.; Shah, S.; Behl, H. S.; Wang, X.; Bubeck, S.; Eldan, R.; Kalai, A. T.; Lee, Y. T.; and Li, Y. 2023. Textbooks Are All You Need. *CoRR*, abs/2306.11644.
- Guo, D.; Lu, S.; Duan, N.; Wang, Y.; Zhou, M.; and Yin, J. 2022. UniXcoder: Unified Cross-Modal Pre-training for Code Representation. In Muresan, S.; Nakov, P.; and Villavicencio, A., eds., *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, 7212–7225. Association for Computational Linguistics.
- Guo, D.; Ren, S.; Lu, S.; Feng, Z.; Tang, D.; Liu, S.; Zhou, L.; Duan, N.; Svyatkovskiy, A.; Fu, S.; Tufano, M.; Deng, S. K.; Clement, C. B.; Drain, D.; Sundaresan, N.; Yin, J.; Jiang, D.; and Zhou, M. 2021. GraphCodeBERT: Pre-training Code Representations with Data Flow. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Guo, D.; Zhu, Q.; Yang, D.; Xie, Z.; Dong, K.; Zhang, W.; Chen, G.; Bi, X.; Wu, Y.; Li, Y. K.; Luo, F.; Xiong, Y.; and Liang, W. 2024. DeepSeek-Coder: When the Large Language Model Meets Programming - The Rise of Code Intelligence. *CoRR*, abs/2401.14196.
- Hoffmann, J.; Borgeaud, S.; Mensch, A.; Buchatskaya, E.; Cai, T.; Rutherford, E.; de Las Casas, D.; Hendricks, L. A.; Welbl, J.; Clark, A.; Hennigan, T.; Noland, E.; Millican, K.; van den Driessche, G.; Damoc, B.; Guy, A.; Osindero, S.; Simonyan, K.; Elsen, E.; Rae, J. W.; Vinyals, O.; and Sifre, L. 2022. Training Compute-Optimal Large Language Models. *CoRR*, abs/2203.15556.
- Huang, J.; Tang, D.; Shou, L.; Gong, M.; Xu, K.; Jiang, D.; Zhou, M.; and Duan, N. 2021. CoSQA: 20, 000+ Web Queries for Code Search and Question Answering. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, 5690–5700. Association for Computational Linguistics.
- Husain, H.; Wu, H.; Gazit, T.; Allamanis, M.; and Brockschmidt, M. 2019. CodeSearchNet Challenge: Evaluating the State of Semantic Code Search. *CoRR*, abs/1909.09436.
- Jiang, X.; Zheng, Z.; Lyu, C.; Li, L.; and Lyu, L. 2021. TreeBERT: A tree-based pre-trained model for programming language. In de Campos, C. P.; Maathuis, M. H.; and Quaghebeur, E., eds., *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI 2021, Virtual Event, 27-30 July 2021*, volume 161 of *Proceedings of Machine Learning Research*, 54–63. AUAI Press.
- Kocetkov, D.; Li, R.; Allal, L. B.; Li, J.; Mou, C.; Ferrandis, C. M.; Jernite, Y.; Mitchell, M.; Hughes, S.; Wolf, T.; Bahdanau, D.; von Werra, L.; and de Vries, H. 2022. The Stack: 3 TB of permissively licensed source code. *CoRR*, abs/2211.15533.
- Li, R.; Allal, L. B.; Zi, Y.; Muennighoff, N.; Kocetkov, D.; Mou, C.; Marone, M.; Akiki, C.; Li, J.; Chim, J.; Liu, Q.; Zheltonozhskii, E.; Zhuo, T. Y.; Wang, T.; Dehaene, O.; Davaadorj, M.; Lamy-Poirier, J.; Monteiro, J.; Shliazhko, O.; Gontier, N.; Meade, N.; Zebaze, A.; Yee, M.; Umapathi, L. K.; Zhu, J.; Lipkin, B.; Oblokulov, M.; Wang, Z.; V. R. M.; Stillerman, J.; Patel, S. S.; Abulkhanov, D.;

- Zocca, M.; Dey, M.; Zhang, Z.; Moustafa-Fahmy, N.; Bhattacharyya, U.; Yu, W.; Singh, S.; Luccioni, S.; Villegas, P.; Kunakov, M.; Zhdanov, F.; Romero, M.; Lee, T.; Timor, N.; Ding, J.; Schlesinger, C.; Schoelkopf, H.; Ebert, J.; Dao, T.; Mishra, M.; Gu, A.; Robinson, J.; Anderson, C. J.; Dolan-Gavitt, B.; Contractor, D.; Reddy, S.; Fried, D.; Bahdanau, D.; Jernite, Y.; Ferrandis, C. M.; Hughes, S.; Wolf, T.; Guha, A.; von Werra, L.; and de Vries, H. 2023. StarCoder: may the source be with you! *CoRR*, abs/2305.06161.
- Li, X.; Gong, Y.; Shen, Y.; Qiu, X.; Zhang, H.; Yao, B.; Qi, W.; Jiang, D.; Chen, W.; and Duan, N. 2022. CodeRetriever: A Large Scale Contrastive Pre-Training Method for Code Search. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, 2898–2910. Association for Computational Linguistics.
- Lin, J.; Dong, H.; Xie, Y.; and Zhang, L. 2024. Scaling Laws Behind Code Understanding Model. *CoRR*, abs/2402.12813.
- Lozhkov, A.; Li, R.; Allal, L. B.; Cassano, F.; Lamy-Poirier, J.; Tazi, N.; Tang, A.; Pykhtar, D.; Liu, J.; Wei, Y.; Liu, T.; Tian, M.; Kocetkov, D.; Zucker, A.; Belkada, Y.; Wang, Z.; Liu, Q.; Abulkhanov, D.; Paul, I.; Li, Z.; Li, W.; Risdal, M.; Li, J.; Zhu, J.; Zhuo, T. Y.; Zheltonozhskii, E.; Dade, N. O. O.; Yu, W.; Krauß, L.; Jain, N.; Su, Y.; He, X.; Dey, M.; Abati, E.; Chai, Y.; Muennighoff, N.; Tang, X.; Oblokulov, M.; Akiki, C.; Marone, M.; Mou, C.; Mishra, M.; Gu, A.; Hui, B.; Dao, T.; Zebaze, A.; Dehaene, O.; Patry, N.; Xu, C.; McAuley, J. J.; Hu, H.; Scholak, T.; Paquet, S.; Robinson, J.; Anderson, C. J.; Chapados, N.; and et al. 2024. StarCoder 2 and The Stack v2: The Next Generation. *CoRR*, abs/2402.19173.
- Lu, S.; Guo, D.; Ren, S.; Huang, J.; Svyatkovskiy, A.; Blanco, A.; Clement, C. B.; Drain, D.; Jiang, D.; Tang, D.; Li, G.; Zhou, L.; Shou, L.; Zhou, L.; Tufano, M.; Gong, M.; Zhou, M.; Duan, N.; Sundaresan, N.; Deng, S. K.; Fu, S.; and Liu, S. 2021. CodeXGLUE: A Machine Learning Benchmark Dataset for Code Understanding and Generation. In Vanschoren, J.; and Yeung, S., eds., *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.
- Mou, L.; Li, G.; Zhang, L.; Wang, T.; and Jin, Z. 2016. Convolutional Neural Networks over Tree Structures for Programming Language Processing. In Schuurmans, D.; and Wellman, M. P., eds., *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, 1287–1293. AAAI Press.
- Nijkamp, E.; Pang, B.; Hayashi, H.; Tu, L.; Wang, H.; Zhou, Y.; Savarese, S.; and Xiong, C. 2023. CodeGen: An Open Large Language Model for Code with Multi-Turn Program Synthesis. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Rozière, B.; Gehring, J.; Gloeckle, F.; Sootla, S.; Gat, I.; Tan, X. E.; Adi, Y.; Liu, J.; Remez, T.; Rapin, J.; Kozhevnikov, A.; Evtimov, I.; Bitton, J.; Bhatt, M.; Canton-Ferrer, C.; Grattafiori, A.; Xiong, W.; Défossez, A.; Copet, J.; Azhar, F.; Touvron, H.; Martin, L.; Usunier, N.; Scialom, T.; and Synnaeve, G. 2023. Code Llama: Open Foundation Models for Code. *CoRR*, abs/2308.12950.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Wang, X.; Wang, Y.; Mi, F.; Zhou, P.; Wan, Y.; Liu, X.; Li, L.; Wu, H.; Liu, J.; and Jiang, X. 2021. Syncobert: Syntax-guided multi-modal contrastive pre-training for code representation. *arXiv preprint arXiv:2108.04556*.
- Wang, X.; Wang, Y.; Wan, Y.; Wang, J.; Zhou, P.; Li, L.; Wu, H.; and Liu, J. 2022. CODE-MVP: Learning to Represent Source Code from Multiple Views with Contrastive Pre-Training. In Carpuat, M.; de Marneffe, M.; and Ruíz, I. V. M., eds., *Findings of the Association for Computational Linguistics: NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, 1066–1077. Association for Computational Linguistics.
- Wang, Y.; Guo, L.; Shi, E.; Chen, W.; Chen, J.; Zhong, W.; Wang, M.; Li, H.; Zhang, H.; Lyu, Z.; and Zheng, Z. 2023. You Augment Me: Exploring ChatGPT-based Data Augmentation for Semantic Code Search. In *IEEE International Conference on Software Maintenance and Evolution, ICSME 2023, Bogotá, Colombia, October 1-6, 2023*, 14–25. IEEE.
- Yu, Z.; Zhang, X.; Shang, N.; Huang, Y.; Xu, C.; Zhao, Y.; Hu, W.; and Yin, Q. 2024. WaveCoder: Widespread And Versatile Enhancement For Code Large Language Models By Instruction Tuning. In Ku, L.; Martins, A.; and Sriku-mar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, 5140–5153. Association for Computational Linguistics.
- Zhang, D.; Ahmad, W.; Tan, M.; Ding, H.; Nallapati, R.; Roth, D.; Ma, X.; and Xiang, B. 2024. CodeSage: Code Representation Learning At Scale. In *The Twelfth International Conference on Learning Representations*.
- Zhang, J.; Lan, Z.; and He, J. 2023. Contrastive Learning of Sentence Embeddings from Scratch. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, 3916–3932. Association for Computational Linguistics.
- Zhao, H.; Hui, J.; Howland, J.; Nguyen, N.; Zuo, S.; Hu, A.; Choquette-Choo, C. A.; Shen, J.; Kelley, J.; Bansal, K.; Vilnis, L.; Wirth, M.; Michel, P.; Choy, P.; Joshi, P.; Kumar, R.; Hashmi, S.; Agrawal, S.; Gong, Z.; Fine, J.; Warkentin, T.; Hartman, A. J.; Ni, B.; Korevec, K.; Schaefer, K.; and Huffman, S. 2024. CodeGemma: Open Code Models Based on Gemma. *CoRR*, abs/2406.11409.