

XSearch: Explainable Code Search via Concept-to-Code Alignment

YIMING LIU*, Shanghai Jiao Tong University, China and Shanghai Innovation Institute, China

RUOFAN LIU*, National University of Singapore, Singapore

YUN LIN†, Shanghai Jiao Tong University, China

ZICONG ZHANG, Shanghai Jiao Tong University, China

WEIYU KONG, Shanghai Jiao Tong University, China

PENGNIAN QI, Huawei Technologies Co., Ltd, China

XIAO CHENG, Huawei Technologies Co., Ltd, China

WEINAN ZHANG, Shanghai Jiao Tong University, China and Shanghai Innovation Institute, China

QIANXIANG WANG, Huawei Technologies Co., Ltd, China

LINPENG HUANG, Shanghai Jiao Tong University, China

With the emergence of deep learning, semantic code search has been widely adopted in both academia and industry. These approaches embed natural-language queries and code snippets into a shared embedding space and retrieve results based on vector similarity. Despite their strong performance on benchmark datasets, they often suffer from poor explainability and generalization. Retrieved code may appear semantically similar yet miss critical functional requirements of the query, while providing no explanation of why the result was retrieved. Moreover, such failures become more severe under distribution shift, where models struggle to generalize to unseen benchmarks.

In this work, we propose *XSearch*, an intrinsically explainable code search framework. Our key insight is that, by relying on global embedding similarity, all existing retrievers inherently take an inductive view. They learn statistical patterns, rather than truly understand the query's functional requirements. Therefore, we address the problem by reformulating code search as a deductive concept alignment problem. At a high level, *XSearch* (i) identifies functional concepts in the query and (ii) explicitly aligns them with corresponding code statements. This explain-then-predict design not only produces inherent concept-level explanations, but also mitigates shortcut learning that harms out-of-distribution generalization. We train an encoder with explicit concept-alignment objectives and perform retrieval through explicit matching between query concepts and code statements. Experiments show that, when trained on CodeSearchNet with a small model size

*Both authors contributed equally to this research.

†Corresponding author.

Authors' Contact Information: Yiming Liu, Shanghai Jiao Tong University, Shanghai, China and Shanghai Innovation Institute, Shanghai, China, liu_yiming@sjtu.edu.cn; Ruofan Liu, National University of Singapore, Singapore, liu.ruofan16@u.nus.edu; Yun Lin, Shanghai Jiao Tong University, Shanghai, China, lin_yun@sjtu.edu.cn; Zicong Zhang, Shanghai Jiao Tong University, Shanghai, China, zhangzico@sjtu.edu.cn; Weiyu Kong, Shanghai Jiao Tong University, Shanghai, China, kwy160034@sjtu.edu.cn; Pengnian Qi, Huawei Technologies Co., Ltd, Shenzhen, China, qipengnian@huawei.com; Xiao Cheng, Huawei Technologies Co., Ltd, Shenzhen, China, chengxiao5@huawei.com; Weinan Zhang, Shanghai Jiao Tong University, Shanghai, China and Shanghai Innovation Institute, Shanghai, China, wnzhang@sjtu.edu.cn; Qianxiang Wang, Huawei Technologies Co., Ltd, Shenzhen, China, wangqianxiang@huawei.com; Linpeng Huang, Shanghai Jiao Tong University, Shanghai, China, lphuau@sjtu.edu.cn.

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

Conference acronym 'XX, Woodstock, NY

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/2018/06

<https://doi.org/XXXXXXXX.XXXXXXX>

(GraphCodeBERT with 125M parameters), *XSearch* improves performance on out-of-distribution benchmarks from 0.02 to 0.33 (15×) over eight state-of-the-art retrievers, and consistently outperforms both encoder-based and decoder-based baselines with up to 7B parameters. A controlled user study further demonstrates that concept-alignment explanations enable users to accept or reject retrieved results both faster and more accurately.

CCS Concepts: • **Software and its engineering** → **Search-based software engineering**; • **Computing methodologies** → **Information extraction**.

Additional Key Words and Phrases: Explainable Code Search, AI for Analysis and Testing

ACM Reference Format:

Yiming Liu, Ruofan Liu, Yun Lin, Zicong Zhang, Weiyu Kong, Pengnian Qi, Xiao Cheng, Weinan Zhang, Qianxiang Wang, and Linpeng Huang. 2026. XSearch: Explainable Code Search via Concept-to-Code Alignment. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 24 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Code search is a fundamental activity in modern software development, consuming up to 50% of the developers' time during programming tasks such as code reuse, debugging, and feature implementation [63, 86]. Modern code search solutions have been shifted from keyword-based lexical matching techniques [5, 27, 52, 53, 55, 58, 64] to semantic retrieval models [16, 21–23, 32, 39, 67]. By mapping queries and code into a shared dense vector space, these models capture high-level semantic similarities. However, despite their success on in-distribution benchmark datasets, state-of-the-art retrievers remain black boxes, and often fail in real-world, out-of-distribution scenarios. We identify two fundamental challenges in existing code search frameworks.

C1: The Interpretability Gap (Post-hoc v.s. Intrinsic). Developers may encounter retrieved code that appears lexically similar yet fails to satisfy some critical requirements of the query, while the model provides no rationale for why the result was retrieved. Given a query–code pair, developers are presented with only a similarity score, without knowing *which part of the code satisfies which part of the query*. This lack of transparency makes it difficult to decide whether to accept or reject the recommended code. While post-hoc explanation techniques (e.g., Attention maps, XCos [81], EXS [70]) exist, they follow a "predict-then-explain" paradigm. Code is first retrieved using a standard retriever, and explanations are generated afterwards using external mechanisms such as knowledge graphs. As a result, the explanation does not influence how representations are learned and how retrieval decisions are made.

C2: The Generalization Collapse. Existing retrievers often fail to generalize when the test data distribution differs from the training data [10, 26]. As illustrated in Figure 1, models trained on CodeSearchNet tend to rely on "low-hanging" cues such as leading tokens (e.g., function name) to match queries with code. While such cues are highly predictive in-distribution, they become unreliable when code structure or naming conventions change. Consequently, retrieval performance collapses in out-of-distribution settings. Recent studies further suggest that increasing model scale to billions of parameters can partially improve robustness, but does not fundamentally eliminate shortcut reliance [10].

Root Cause: Inductive Matching vs. Deductive Validation. We argue that both challenges stem from a fundamental inductive bias in existing retrievers. Current models are trained to *inductively learn* statistical correlations between queries and code via global embedding matching with binary supervision. However, developers assess code *deductively* by verifying whether each functional requirement in the query is satisfied. This mismatch explains why scaling model capacity alone cannot eliminate shortcut learning: the problem lies not in model size, but in the learning paradigm itself.

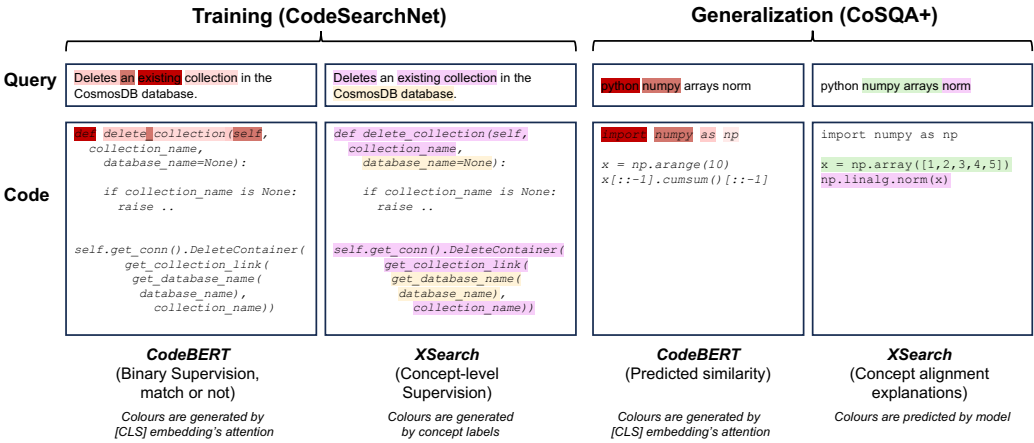


Fig. 1. A CodeBERT retriever trained on CodeSearchNet exhibits shortcut learning, relying on leading tokens (e.g., function names) and failing on CoSQA+, where key semantics appear elsewhere. In contrast, XSearch retrieves by concept-to-code alignment, making the decision traceable and more robust. (Colours: CodeBERT uses post-hoc [CLS] attention. XSearch shows predicted concept highlights and alignments.)

Motivated by this observation, we propose XSearch, which shifts code search from inductive matching to deductive alignment. Concretely, we reformulate retrieval as a *concept-to-code alignment problem*, where the model must explicitly verify that each query concept is grounded in the candidate code. Our key insight is that code search queries are often *compositional*, consisting of multiple functional concepts (e.g., actions, entities, and constraints). The correct code typically implements these concepts in certain localized code regions. Instead of collapsing the entire query and code into single vectors, XSearch (i) first decomposes a query into concepts, (ii) identifies salient code spans that implement similar functionality, and (iii) then computes an optimal concept-to-code alignment to rank candidates and generate explanations. Importantly, concept alignment is integrated into both training and retrieval, encouraging the model to cover multiple functional aspects of the query rather than relying on global similarity alone. This design produces concept-level explanations that can help developers assess whether retrieved code satisfies their functional intent.

We evaluate XSearch against six state-of-the-art encoder-only retrievers (CodeBERT [16], UniX-Coder [22], GraphCodeBERT [23], CoCoSoDa [69], HedgeCode [6], RAPID [14]) as well as two decoder-only retrievers (Qwen2.5-Coder-7B-lt-SupCon-CSN and CodeLlama-lt-SupCon-CSN [10]). Our model is fine-tuned on the GraphCodeBERT encoder with approximately 125M parameters. On out-of-distribution benchmarks, existing retrievers struggle significantly, even for decoder models with 7B parameters. XSearch effectively bridges this gap, improving performance from negligible levels to practical usability (e.g., boosting MRR from 0.02 to 0.33). In addition, a controlled user study demonstrates that concept-alignment explanations enable participants to accept or reject retrieved results 38% faster and 10.5% more accurately than when only similarity scores are provided.

We summarize the contributions of this paper as follows:

- We identify a fundamental limitation of existing code retrievers: their inductive formulation encourages shortcut learning. We address this by reformulating code search as deductive concept alignment, enabling the model to verify each functional requirement explicitly.
- We propose XSearch, an intrinsically explainable code retrieval framework, where concept-level alignments are unified into the model’s prediction, rather than generated in a post-hoc way.

- We develop an LLM-assisted label augmentation pipeline to construct concept annotations for training an alignment-aware retrieval model. We open-source both the label augmentation framework and the dataset in [2].
- We extensively evaluate *XSearch* with eight state-of-the-art code retrievers on seven code search tasks. In addition, we design a user study consisting of 20 participants over 20 code search tasks. The results demonstrate great improvements in both explainability and performance generalizability.

2 Preliminary

In this section, we first introduce related work on code search, then we present our motivating example.

2.1 Background

LLMs for Code. In recent years, large language models (LLMs) have been widely adopted for code-related tasks. Some serve as foundational models pre-trained on general code-solving tasks, while others are specialized for specific domains. General-purpose code LLMs are pre-trained on typical NLP tasks such as masked language modeling [22, 84, 85, 91], span denoising [22, 84, 91], next-sequence prediction [22, 24], and fill-in-the-middle prediction [24, 44]. Some research works incorporate code-specific tasks such as AST edge prediction [83], data flow edge prediction [23], identifier tagging [83, 85], and code-NL contrastive learning [22, 83–85, 91].

Meanwhile, specialized LLMs are designed to address specific tasks. The tasks can be grouped into the following categories: 1) NL-to-code generation: code generation from documentation [4, 8, 11, 28, 45], 2) NL-to-code retrieval: code search [6, 33, 46, 57, 69, 89]. 3) Code-to-NL generation: code summarization [1, 68, 74, 78] and code-to-comment generation [3, 38, 43]. 4) Code-to-code generation: code completion [35, 51, 90], code repair [15, 17, 31, 62, 88] and code translation [9, 29, 40, 66, 75, 79]. Beyond these applications, LLMs are increasingly used to streamline the software development process, including tasks such as commit message generation [13, 37, 42, 87], automated code review [77, 80], refactoring suggestions [73], and detecting security vulnerabilities [25, 34, 47, 48, 59, 72, 82].

Code Search. In this work, we focus on the code search task. Code search aims to retrieve relevant code snippets given a natural-language query [7, 12, 39, 50, 76]. It can be integrated into downstream applications such as code generation [7], fault localization [36], and program repair [92]. Existing approaches can be broadly categorized into *keyword-based* methods and *embedding-based* semantic methods.

- **Keyword-based code search.** Early code search systems relied on lexical matching between a natural-language query and code artifacts, using information-retrieval techniques over identifiers, comments, and API names [5, 53, 55]. To reduce vocabulary mismatch, follow-up work explored query expansion and reformulation, e.g., by adding synonyms or related terms to the query [27, 52, 58, 64]. While efficient and transparent, keyword-based retrievers are limited by surface-form overlap and often fail when relevant code uses different wording, naming conventions, or abstractions.
- **Embedding-based semantic code search.** To overcome lexical limitations, embedding-based approaches map queries and code snippets into a shared embedding space, where relevance is measured by vector similarity. Early neural network approaches such as Deep Code Search [21] and CodeSearchNet [32] learn joint representations for code and natural language. Recent Transformer-based models, including CodeBERT [16], GraphCodeBERT [23], UniXcoder [22], and CodeRetriever [46], leverage large-scale pretraining and self-attention to further improve

retrieval accuracy. More recent methods such as CoCoSoDa [69] and HedgeCode [6] adopt contrastive learning objectives to strengthen representation alignment. Besides, recent efforts are being made to improve accuracy by deepening search steps [19], improving searching efficiency [20], and enhancing data augmentation strategies [41].

2.2 Motivating Example

Figure 1 shows an observed shortcut learning behavior for an encoder-based code retriever trained with CodeBERT [16] on the CodeSearchNet benchmark [32]. The figure shows four query-code pairs. The first pair is a training example from CodeSearchNet with a binary label (i.e., whether they are relevant or not). The second pair is the same training example with our proposed concept-level label augmentation, where query concepts are explicitly aligned with code units. The third pair is a top-ranked result retrieved by CodeBERT, on out-of-distribution benchmark CoSQA+ [18]. And the fourth pair is the top-ranked result retrieved by *XSearch*. In addition, we highlight the focal tokens that contribute the most to the model’s prediction. For CodeBERT, contributions are measured by post-hoc methods, and we trace the attention weights by using [CLS] embedding as query and all tokens as keys. For *XSearch*, we visualize the predicted concept-bearing tokens and their concept alignments directly.

Observation: Shortcut Learning in CodeBERT. On CodeSearchNet, CodeBERT assigns high contribution to leading tokens such as function names. This behavior is effective in-distribution (first pair), since many code sample conveys their functionality in function names. However, when applied to out-of-distribution data where such naming conventions do not hold, the model relies on the same positional cues and retrieves incorrect results, as shown in the third pair.

Key Difference with Concept-level Supervision. In contrast, *XSearch* explicitly models retrieval as a concept alignment problem. Instead of relying on global embedding similarity, it aligns each functional concept in the query with corresponding code units. As shown in the fourth pair, the retrieve code covers all the required concepts “numpy arrays” and “norm”. This reduces the model’s tendency to over-reliance on positional or naming cues.

Our Solution: Label Augmentation. Different from the traditional data augmentation to improve the model generalizability, we adopt a novel label augmentation in this work. Specifically, we annotate the training samples by (1) how to extract the concept in the query and (2) how those concepts are present and absent in the code. This mirrors how developers assess code relevance in practice, where missing any required operation often renders a candidate unusable. By training the model to reason over concept coverage instead of surface patterns, *XSearch* learns representations that are both more interpretable and more robust to distribution shift.

3 Approach

3.1 Overview

The input is a natural-language query n and the goal is to retrieve code snippets $c \in C$ that best satisfy the query. We propose a novel retrieval model architecture, as shown in Figure 2. In addition to retrieving relevant code, *XSearch* also produces explanations by explicitly aligning code lines with the concepts mentioned in the query. We develop *XSearch* through three stages:

- **Stage 1: Concept Label Augmentation (Section 3.2).** We annotate essential concepts in queries and identify the corresponding code units that implement each concept. This stage produces token-level query-code alignment annotations.
- **Stage 2: Alignment-Aware Model Training (Section 3.3).** Using the annotated alignments, we train an alignment-aware retrieval model that jointly identifies concept-bearing tokens and aligns semantically related query-code pairs in the embedding space.

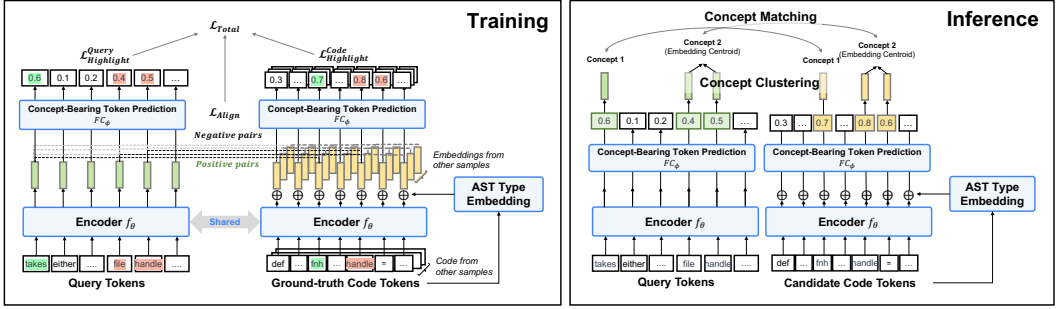


Fig. 2. *XSearch* Model Architecture. During training (LHS), concept-aligned query and code token spans form positive pairs, while non-aligned spans form negative pairs. A shared encoder maps query and code tokens to contextual embeddings. For code tokens, AST type embeddings are added to incorporate structural information. A linear probing head predicts concept-bearing tokens at the token level, and an alignment objective trains the model to associate query concepts with their corresponding code spans. At inference time (RHS), the model identifies concept-bearing tokens in both query and code, aggregates them into concept representations, and matches query concepts to code concepts based on embedding similarity.

Table 1. Concept Components and Examples.

Component	Description
Action	Operation to perform (<i>delete, return</i>)
Entity	Target object (<i>collection, file handle</i>)
Modifier	Constraint/attribute (<i>existing, open</i>)
Concept Example	Composition
<i>delete existing collection</i>	Action + Modifier + Entity
<i>open file handle</i>	Modifier + Entity

- **Stage 3: Explainable Query-to-Code Retrieval (Section 3.4).** At inference time, the model aligns predicted concept-bearing tokens between the query and candidate code snippets, retrieves the top code by similarity, and outputs the corresponding concept alignment maps as explanations.

3.2 Concept Label Augmentation

We perform *label augmentation* rather than data augmentation. For each existing query-code pair, we annotate (i) concept spans in the query and (ii) alignments from each concept to the code units that implement it. This produces fine-grained supervision without creating new training examples. Because such annotations must be both scalable and reliable, we combine LLM-based annotation with strict validation.

3.2.1 Concept Definition. We define a concept as a set of query tokens that together express a single, indivisible functional requirement. Internally, a concept is compositional: it typically consists of an **Action** or **Entity**, and optionally combined with **Modifiers** that specify constraints or attributes (Table 1). For example, “delete existing collection” forms a unified concept with an action “delete”, an entity “collection”, and a modifier “existing”. Each query token is assigned to at most one concept (or marked as non-concept). We do not assume concepts to be perfectly defined. Rather, we only require that constraint-bearing units can be identified with reasonable consistency under our definition.




Fig. 3. LLM-assisted Annotation Pipeline.

3.2.2 *LLM-Assisted Annotation Pipeline.* Given a query n and its paired code snippet c in the dataset, we annotate them in three steps (Figure 3).

- (1) **Step 1: Query Concept Decomposition.** We ask the LLM to identify the concepts in the query. A concept a is represented as a set of query tokens $\mathcal{N}_a = \{n_i \mid \phi(n_i) = a\}$, where $\phi(\cdot)$ assigns each token to a concept or to *None*.
- (2) **Step 2: Code Summarization.** We ask the LLM to split the code into code units (by default, code lines, consecutive lines can be merged when they form one coherent statement), and ask it to write a short natural-language description d for each unit. This provides an intermediate textual description that is easier to align with query concepts.
- (3) **Step 3: Code Alignment.** We align each query concept a to one or more code units that implement it. We then map aligned units back to code tokens, producing $C_a = \{c_j \mid \phi(c_j) = a\}$. Unlike queries, a code unit (and its tokens) can be linked to multiple concepts.

3.2.3 *Assertion-Based Validation.* To improve annotation reliability, we validate each LLM output using two types of assertions: (1) **format assertions** to ensure that the output can be parsed, and (2) **consistency assertions** to ensure that every concept span and aligned code unit exactly appears in the original query and code. If any assertion fails, we return the error messages to the LLM and retry from Step 1. We discard samples that fail after R retries. In our data construction, 72.98% pass with no retry, 97.11% pass within two retries, and all accepted samples pass within five retries.

3.2.4 *Small-Scale Manual Verification.* To assess annotation quality, we randomly sampled 1,000 annotated pairs per programming language and asked two developers (with 3 years+ programming experience) to independently judge whether (i) query concept spans follow the definition in Table 1, and (ii) aligned code units reasonably correspond to the intended concepts. The inter-annotator agreement is high ($\kappa = 0.9$), suggesting that the annotations are largely consistent under our definition, although we do not claim this procedure fully eliminates semantic errors.

3.3 Alignment-Aware Model Training

To jointly learn (i) concept-bearing tokens and (ii) concept-to-code alignments, we optimize two objectives: a token-level **highlight loss** and a span-level **alignment loss**. The highlight loss predicts whether each token should be highlighted as concept-bearing. The alignment loss pulls each query concept span close to its aligned code span while pushing it away from hard negatives.

3.3.1 **Highlight Loss.** We add a lightweight probing head on top of the encoder to predict a highlighting probability for each token. Given an input sequence x (a query or code snippet), the encoder f_θ outputs contextual hidden states $h_{i=1}^N$, where $h_i = f_\theta(x_i)$ is the representation of token x_i . A linear classifier produces $p_i \in (0, 1)$:

$$p_i = \sigma(\text{FC}_\phi(h_i)) = \sigma(\text{FC}_\phi(f_\theta(x_i))), \quad x_i \in \{\text{query tokens, code tokens}\}. \quad (1)$$

Since concept-bearing tokens are sparse, we use focal loss to address the class imbalance problem [49]:

$$\mathcal{L}_{\text{high}} = - \sum_{i=1}^N \left[\alpha (1 - p_i)^\gamma y_i \log p_i + (1 - \alpha) p_i^\gamma (1 - y_i) \log(1 - p_i) \right], \quad (2)$$

where $y_i \in \{0, 1\}$ indicates whether token x_i is concept-bearing, and α, γ are focal-loss hyperparameters. Note that when $\gamma = 0$ and $\alpha = 0.5$, the focal loss reduces to standard binary cross-entropy.

For code tokens, we additionally incorporate coarse-grained AST node types (20 types) via a type embedding layer. Type embeddings are added to token representations before the probing head. We use separate probing heads for queries and code.

3.3.2 Alignment Loss with Hard Negative Sampling. Let a query concept span be \mathcal{N}_a (a set of query tokens belonging to concept a), and its aligned code span be \mathcal{C}_a (a set of code tokens aligned to the same concept). We represent a span by mean-pooling its token embeddings:

$$f(\mathcal{S}) = \frac{1}{|\mathcal{S}|} \sum_{x_i \in \mathcal{S}} f_\theta(x_i), \quad \mathcal{S} = \mathcal{N}_a \text{ or } \mathcal{C}_a. \quad (3)$$

We then apply a contrastive objective: for each \mathcal{N}_a , treat $(\mathcal{N}_a, \mathcal{C}_a)$ as a positive pair and sample K hard negatives $\mathcal{S}_{a,[K]}^-$ from (i) *intra-sample negatives* (other concepts in the same query-code pair) and (ii) *inter-sample negatives* (concept spans from other pairs in the batch). The full algorithm is presented in Algorithm 1.

Specifically, to prioritize “hard” negatives that are spuriously similar under the current model, we score each negative \mathcal{C}_b ($b \neq a$) by:

$$\text{score}_b = \underbrace{\cos(f_\theta(\mathcal{N}_a), f_\theta(\mathcal{C}_b))}_{\text{current model}} - \underbrace{\cos(g(\mathcal{N}_a), g(d_b))}_{\text{reference text model}}, \quad (4)$$

where $f_\theta(\cdot)$ uses our encoder representations, and $g(\cdot)$ is a frozen sentence encoder (all-mpnet-base-v2 [71]) pretrained on semantic similarity tasks. The reference model is used *only* to estimate negative hardness and is independent of our model. Its input is \mathcal{N}_a and d_b , the natural-language description of the code span \mathcal{C}_b (from Step 2 in the annotation pipeline described in Section 3.2). Intuitively, a high score indicates that the current model considers $(\mathcal{N}_a, \mathcal{C}_b)$ similar, while the reference text model considers \mathcal{N}_a and d_b semantically mismatched. In other words, we prioritize negatives that the current model mistakenly finds similar, these are the most informative for training.

After negative sampling, we employ the contrastive learning loss. We adopt InfoNCE loss [60] as follows.

$$\mathcal{L}_{\text{Align}} = \sum_{\mathcal{N}_a} -\log \frac{\exp(u_a/\tau)}{\exp(u_a/\tau) + \sum_{\mathcal{C}_b \in \mathcal{S}_{a,[K]}^-} \exp(v_{a,b}/\tau)}, \quad (5)$$

$$\text{where } u_a = \cos(f_\theta(\mathcal{N}_a), f_\theta(\mathcal{C}_a)), v_{a,b} = \cos(f_\theta(\mathcal{N}_a), f_\theta(\mathcal{C}_b)).$$

Here, \mathcal{N}_a is the a -th query-concept span’s embedding, \mathcal{C}_a is its positive code-span embedding, $\mathcal{S}_{a,[K]}^-$ is the set of K hardest negatives for \mathcal{N}_a , and τ is the temperature hyper-parameter. The overall training objective becomes:

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{High}}^{\text{Code}} + \mathcal{L}_{\text{High}}^{\text{query}} + \mathcal{L}_{\text{Align}} \quad (6)$$

Algorithm 1 Hard Negative Sampling for Alignment Loss

Require: A batch of query-code pairs; Current model f_θ , reference model f_{ref} ; Number of negatives per concept as K

Ensure: Query concepts and their Top K hard negatives: $\{(\mathcal{N}_a, \mathcal{S}_{a,[K]}^-)\}$

```

1: for query concept  $\mathcal{N}_a$  do
2:    $\mathcal{S}_a^- \leftarrow \{C_b \mid b \neq a\}$  // Collect all negative code spans
3:   for code span  $C_b \in \mathcal{S}_a^-$  do
4:      $score_b \leftarrow \cos(f_\theta(\mathcal{N}_a), f_\theta(C_b)) - \cos(g(\mathcal{N}_a), g(d_b))$ 
5:   end for
6:   // Select Top  $K$  hardest negatives
7:   Sort  $\mathcal{S}_a^-$  by  $score_b$  in descending order
8:    $\mathcal{S}_{a,[K]}^- \leftarrow$  first  $K$  elements of  $\mathcal{S}_a^-$ 
9: end for
10: return  $\{(\mathcal{N}_a, \mathcal{S}_{a,[K]}^-)\}$ 

```

We use equal weights as the three losses are empirically on comparable scales.

3.4 Explainable Query-to-Code Retrieval

At inference time, the input consists of a natural-language query N_{query} and a reference codebase C_{ref} .

Concept Clustering for Query. Given a query, we run inference with the retrieval model and obtain the highlight score p_i for token i . Tokens with $p_i > \delta_{highlight}$ are selected as concept-bearing and grouped using agglomerative hierarchical clustering with cosine similarity. Clusters whose centroid similarity exceeds a threshold $\delta_{cluster}$ are merged. Each resulting cluster C is represented by a centroid

$$\mathbf{h}_C = \frac{1}{|C|} \sum_{n_i \in C} f_\theta(n_i), \quad (7)$$

which serves as a concept-level representation of the query. We set $\delta_{highlight} = 0.4$ and $\delta_{cluster} = 0.8$ (sensitivity analysis in Figure 4).

Concept Matching for Code. For code, we adopt a line-level representation rather than clustering. Unlike queries where concept boundaries must be inferred, code has explicit structural boundaries: each line typically corresponds to a single statement. For each non-empty code line l , we check whether it contains at least one concept-bearing token (i.e., $\exists c_j \in l$ such that $p_j > \delta_{highlight}$). Lines without any concept-bearing tokens are excluded from matching. For retained lines, let C_l denote *all* tokens in line l , and we compute a line-level embedding:

$$\mathbf{e}_l = \frac{1}{|C_l|} \sum_{c_j \in C_l} f_\theta(c_j). \quad (8)$$

Given query-concept centroids $\{\mathbf{h}_{C_k}\}_{k=1}^K$ and code-line embeddings $\{\mathbf{e}_l\}_{l=1}^L$, we apply greedy matching: each query concept is matched to the code line with the highest cosine similarity. The overall similarity is the average across all matches:

$$\text{sim}(q, c) = \frac{1}{K} \sum_{k=1}^K \max_l \cos(\mathbf{h}_{C_k}, \mathbf{e}_l). \quad (9)$$




Fig. 4. Hyperparameter sweep on the CodeSearchNet-python [32] test set. We vary $\delta_{highlight}$ and $\delta_{cluster}$ within a focused range and report MRR. Darker color indicates higher MRR.

This ensures that the retrieved code must address all query concepts to score highly. Note that *XSearch* is designed to be conservative at inference time, prioritizing constraint satisfaction over partial matches.

3.4.1 Output Explanation Format. Figure 9b shows an example output. Given the query, *XSearch* identifies two concepts and returns the top-2 code snippets. For each snippet, it highlights the lines that support each concept. The top-2 results both implement dictionary merging and therefore receive high similarity. Code B is closer because it satisfies the requirement on “overwriting existing keys”, while Code A doesn’t. Overall, *XSearch* not only ranks code snippets by semantic relevance, but also provides explicit concept-to-code alignments, reducing developers’ cognitive effort when deciding whether to accept or reject a recommendation. We evaluate this benefit in our user study (Section 4.7).

3.4.2 Complexity Analysis. We analyze the inference complexity of *XSearch*. Let a and b denote the token lengths of the query and code, respectively. The encoder forward pass costs $\mathcal{O}(a)$ and $\mathcal{O}(b)$. Let k be the number of highlighted tokens in the query ($k < a$). Query concept clustering costs $\mathcal{O}(k^2)$, while centroid computation is $\mathcal{O}(k)$. For code, computing line-level embeddings (Equation 8) costs $\mathcal{O}(b)$, and concept matching costs $\mathcal{O}(KL)$, where K is the number of query concepts and L is the number of code lines. Overall, the inference complexity is $\mathcal{O}(a + k^2 + |C_{ref}| \cdot (b + KL))$. Since k , K , and L are small in practice, the complexity is dominated by encoder computation, comparable to a standard bi-encoder.

4 Experiments

4.1 Hyperparameter Setup

In the training stage, we fine-tune GraphCodeBERT [56] on CodeSearchNet [32], use a batch size of 32, an epoch number of 10, Adam optimizer with initial learning rate $2e - 5$, and linear learning rate decay. We set $\gamma = 2.0$, $\alpha = 0.5$ for highlight loss. For alignment loss, we use the reference model as “all-mpnet-base-v2” [71], which is pretrained on natural language sentence pairs. We sample the top-50 ranked negative code for alignment loss computation, and we set temperature $\tau = 0.1$ (same as [23]). At the inference stage, we first identify concept-bearing tokens by thresholding the highlight scores with $\delta_{highlight}$ and then merge the retained tokens into concepts using the cluster threshold $\delta_{cluster}$. Based on a small-scale experiment on the test set (Figure 4), we set $(\delta_{highlight}, \delta_{cluster}) = (0.4, 0.8)$ and find the performance to be stable under modest variations of the thresholds.

4.2 Research Questions

To evaluate our approach, we conduct extensive experiments designed to answer four key research questions (RQs):

RQ1: Overall Retrieval Performance: How does *XSearch* compare with existing methods, on both in-distribution and out-of-distribution code search benchmarks?

RQ2: Faithfulness of Concept Explanations: How accurate are the two explanation components produced by *XSearch*: (i) highlighted concept-bearing tokens and (ii) concept-to-code alignments?

RQ3: Ablation Study: How do the key components of *XSearch* (e.g., concept highlighting, negative sampling in alignment training) affect retrieval performance?

RQ4: User Study: Do concept-alignment explanations help developers judge retrieved results more quickly and more accurately than similarity scores alone?

Given the space limit, more experimental results, demonstrations, and qualitative analysis are available at [2].

4.3 RQ1: Overall Retrieval Performance

4.3.1 Setup. We evaluate *XSearch* on both in-distribution (ID) and out-of-distribution (OOD) benchmarks. CodeSearchNet (CSN) [32] is used for ID evaluation, providing large-scale one-to-one query–code pairs across six languages. CoSQA+ [18] is used for OOD evaluation, where each query is paired with multiple relevant code snippets, reflecting more realistic retrieval settings. All models are fine-tuned on CSN and evaluated on CSN test sets for ID performance, and directly transferred to CoSQA+ for OOD evaluation. Table 2 reports dataset statistics.

Table 2. Datasets Summary

Dataset	Language	#Train	#Valid	#Test	#Codebase
CodeSearchNet	Ruby	24,927	1,400	1,261	4,360
	JavaScript	58,025	3,885	3,291	13,981
	Java	164,923	5,183	10,955	40,347
	Go	167,288	7,325	8,122	28,120
	PHP	241,241	12,982	14,014	52,660
	Python	251,820	13,914	14,918	43,827
CoSQA+	Python	16,481	2,058	2,065	51,516

4.3.2 Baselines. We select the baselines to verify two key hypotheses: (i) whether performance gains on out-of-distribution data can be achieved by introducing concept-to-code alignment objectives while keeping the same encoder architecture, and (ii) whether such generalization can be achieved by scaling up model size alone, without reformulating the retrieval objective.

Encoder-based Code Retrievers. We include six widely used encoder-based retrievers. All selected encoder baselines follow the bi-encoder retrieval paradigm, where queries and code are independently mapped into a shared embedding space. CodeBERT [16], UniXcoder [22], and GraphCodeBERT [23] mainly enhance representations through pretraining strategies (e.g., masked language modeling, cross-modal, and multi-task pretraining). CoCoSoDa [69], HedgeCode [6], and RAPID [14] aim to improve robustness by modifying training objectives or data (e.g., contrastive learning, adversarial examples, or pseudo-labeled samples). These models encode queries and code [CLS] token into a shared embedding space, and perform retrieval by their [CLS] similarity. *XSearch* builds upon the GraphCodeBERT architecture and introduces an additional concept-bearing prediction head (Figure 2).

Table 3. Retrieval Performance Comparison. Best Scores are Shown in **Bold**.

Model	In-Distribution (CodeSearchNet)						Out-of-Distribution (CoSQA ⁺)	
	Ruby (MRR)	JavaScript (MRR)	Go (MRR)	Python (MRR)	Java (MRR)	PHP (MRR)	MRR	MMRR
CodeBERT [16]	0.672	0.621	0.882	0.672	0.677	0.626	0.018	0.012
UniXCoder [22]	0.752	0.683	0.917	0.723	0.724	0.674	0.032	0.022
GraphCodeBERT [23]	0.703	0.644	0.897	0.692	0.691	0.649	0.020	0.010
CoCoSoDa [69]	0.803	0.748	0.923	0.742	0.753	0.681	0.028	0.022
HedgeCode [6]	0.785	0.748	0.883	0.718	0.731	0.686	0.026	0.020
RAPID [14]	0.745	0.767	0.925	0.739	0.726	0.660	0.054	0.036
CodeLlama-7B-It-SupCon-CSN [10]	0.569	0.645	0.747	0.748	0.742	0.686	0.163	0.114
Qwen2.5-Coder-7B-It-SupCon-CSN [10]	0.645	0.656	0.772	0.793	0.754	0.688	0.240	0.178
Ours	0.725	0.654	0.892	0.720	0.690	0.604	0.332	0.242

Decoder-based Code Retrievers. Recent work [10] has shown that decoder-only LLMs achieve better generalization performance than encoder-only LMs when fine-tuned on code retrieval tasks. Since decoder models do not have a special pooling token, they compute sequence-level embedding by taking the average embedding over all query/code tokens, and perform retrieval using vector similarity. Following this setting, we include two strong LLM-based baselines released by [10]: *Qwen2.5-Coder-7B-It-SupCon-CSN* and *CodeLlama-7B-It-SupCon-CSN*. They are fine-tuned from Qwen2.5-Coder-7B [30] and CodeLlama-7B [65] on CodeSearchNet, respectively. These models represent the “scale-up” approach without changing the retrieval objective.

4.3.3 Evaluation Metrics. For each query, the retriever ranks code from the codebase by their relevance (the codebase size is detailed in Table 2). We employ two metrics to evaluate retrieval performance.

- **Mean Reciprocal Rank (MRR).** MRR is the standard metric for ranking-based retrieval tasks [16, 56]. For each query, the reciprocal rank (RR) is defined as the inverse of the rank position of the first relevant result. MRR is computed as the average RR over all queries:

$$\text{MRR} = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \frac{1}{\text{rank}_i} \quad (10)$$

where $|Q|$ is the number of queries, and rank_i is the rank position of the first relevant code snippet for the i -th query.

- **Mean Multi-choice Reciprocal Rank (MMRR).** Unlike CodeSearchNet, queries in CoSQA⁺ are associated with multiple relevant code snippets. Following prior work, MMRR is used to evaluate CoSQA⁺ [10]. It computes the highest rank among all correct answers:

$$\text{MMRR} = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \max_{j \in R_i} \frac{1}{\text{rank}_{i,j}} \quad (11)$$

where R_i denotes the set of relevant code snippets for the i -th query, and $\text{rank}_{i,j}$ is the rank position of the j -th relevant snippet. We adopt MMRR because it better reflects retrieval quality in multi-answer scenarios by evaluating whether a model can consistently rank any relevant result highly, rather than focusing on a single labeled match.

4.3.4 Results. Table 3 reports the retrieval performance of all models on both in-distribution (ID) CodeSearchNet and out-of-distribution (OOD) CoSQA⁺. On the in-distribution (ID) CodeSearchNet

Table 4. Efficiency Comparison.

Metric	GraphCodeBERT [56]	Qwen2.5-Coder-7B-It-SupCon-CSN [10]	XSearch
Parameters	125M	7B	125M
Inference (per query)	0.01s	0.10s	0.03s
GPU Memory	2.1 GB	28.5 GB	2.3 GB

Baseline CoCoSoDa Prediction	XSearch Prediction
<p>Query Angle between two vectors using python</p> <p>Retrieved Top 1 Code (Score = 27.07) ❌ <pre>def angle_between_vectors(vec1, vec2): return tuple([vec1[i] / vec2[i] for i in range(len(vec2))])</pre></p> <p>Ground Truth Code (Rank > 5) <pre>def angle_between_vectors(x, y): dp = dot_product(x, y) if dp == 0: return 0 xm = magnitude(x) ym = magnitude(y) return math.acos(dp / (xm * ym)) * (180 / math.pi)</pre></p>	<p>Query Angle between two vectors using python</p> <p>Retrieved Top 1 Code (Score = 0.88) ✅ <pre>def angle_between_vectors(x, y) // both [] and [] dp = dot_product(x, y) if dp == 0: return 0 xm = magnitude(x) ym = magnitude(y) return math.acos(dp / (xm * ym)) * (180 / math.pi)</pre></p> <p>Retrieved Top 2 Code (Score = 0.83) <pre>def angle(x0, y0, x1, y1): return degrees(atan2(y1 - y0, x1 - x0))</pre></p>

Fig. 5. Failure Case of CoCoSoDa [69].

benchmark, encoder-only models show surprisingly high performance, often outperforming much larger 7B decoder models. For example, CoCoSoDa achieves an MRR of 0.803 on the Ruby dataset. However, their performances are near-zero on CoSQA+. Decoder-based models obtain better transferability to OOD data (e.g., Qwen2.5-Coder-7B-It-SupCon-CSN achieves an MRR of 0.24 on CoSQA+). Nevertheless, these gains come at the cost of substantially larger training effort and higher inference overhead (Table 4). XSearch achieves a better balance between efficiency and generalizability. Its ID performance can be slightly lower when queries contain extra or overly specific constraints not reflected in the paired code, whereas it outperforms all baselines on OOD benchmark (MRR of 0.33). Notably, these improvements are achieved by reformulating the retrieval objective rather than increasing model scale. We analyze the failure cases on in-distribution set in Section 4.4.3.

4.4 Case Studies

4.4.1 Why does XSearch outperform encoder-based baselines? Figure 5 shows a failure case of the strongest encoder-based baseline, CoCoSoDa. The failure reason is similar to that of CodeBERT in Section 2.2. The query asks for computing the *angle* between two vectors. However, CoCoSoDa retrieves a code snippet that performs *division* between two vectors, which is lexically similar to the query as both mention “two vectors”, but is functionally irrelevant. This happens because bi-encoder retrieval optimizes similarity between single [CLS] representations, without enforcing coverage of all query concepts. As a result, representations may be dominated by salient or frequent tokens (e.g., “two vectors”), leading to incorrect matches. In contrast, XSearch ranks the ground-truth implementation as the Top-1 result with a high similarity score. The Top-2 result is also semantically correct, except that the two vectors are expressed in coordinate form. We show more examples on our anonymous website [2].

4.4.2 Why does XSearch outperform decoder-based baselines? Despite the strong reasoning capabilities of large decoder-based models, Figure 6 demonstrates that they can still overlook specific functional constraints when processing OOD queries. In this example, the query requires transforming a string into a *dictionary* object. However, the baseline (Qwen2.5-Coder-7B-It-SupCon-CSN) retrieves code that merely splits and prints the string, failing to construct the required dictionary

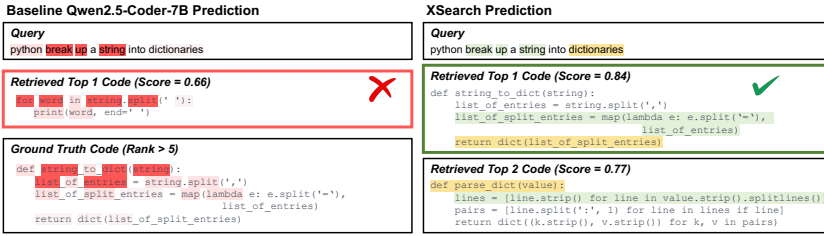


Fig. 6. Failure Case of Qwen2.5-Coder-7B-It-SupCon-CSN [10].

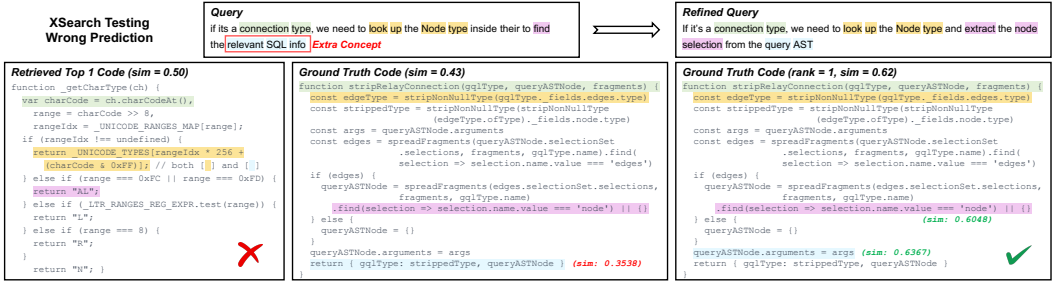


Fig. 7. Failure Case of XSearch. *Left*: Retrieved Top-1 code for original query. *Middle*: Ground-truth code for original query, where the query includes an extra concept (“relevant SQL info”) not reflected in the code. *Right*: Refined query (replacing the extra concept with “relevant query AST”) and the ground-truth code.

structure. This behavior can be attributed to the representation formulation used in decoder-based retrieval, where a sequence is represented by the average over all token embeddings [10]. Under this formulation, the contributions of constraint-bearing tokens can be diluted by other tokens in the sequence. To verify this effect, we compute the embedding shift induced by masking individual query tokens. We observe that masking the core constraint “dictionaries” leads to only a negligible change in the averaged embedding. This indicates that the “dictionaries” influence is suppressed by the pooling operation, and the final representation is dominated by more generic action-related tokens such as “break up”. In contrast, XSearch avoids this dilution effect by computing concept alignments independently, ensuring that every concept directly influences retrieval decisions.

4.4.3 Why does XSearch slightly under-perform on in-distribution testing? To clarify the ID performance gap, we analyzed 93 failure cases and found that the most prevalent error (32.3%) stems from redundant query concepts that lack implementation in the code. Other main factors include single-concept collapse (22.6%), imperfect alignment (15.1%), and missed concepts (11.8%). XSearch adopts a stricter matching criterion than prior retrievers. While most baselines optimize a single global query-code similarity where partial semantic overlap can yield a high score, XSearch emphasizes **functional requirement completeness** by requiring coverage of all concepts in the query. As a result, when a query includes extra or overly specific functional-requirement concepts that are not implemented in the corresponding code, XSearch will intentionally assign a lower similarity score rather than rewarding partial matches. Figure 7 illustrates this leading error mode (the 32.3% category) through a concrete example. The original query includes an extra concept, “relevant SQL info”, which is not reflected in the ground-truth code, indicating that the query and code are not synchronized in their information content. This missing functional requirement concept reduces the overall query-code score, and the top retrieved result is also

low-scored, suggesting limited reliability. For developers, returning a high-scored snippet that fails (or cannot be verified) to satisfy an explicitly stated functional requirement can be more misleading than returning no confident result. When refining the query by replacing the extra concept with a code-grounded description (“extract the node selection from the query AST”), the ground-truth code aligns with all query concepts and achieves a substantially higher similarity. More examples are available on our anonymous website [2].

4.5 RQ2: Faithfulness of Concept Explanations




Fig. 8. Alignment and Highlight Performance Across Languages.

4.5.1 Setup. We define two metrics for assessing the accuracy of two explanation components produced by XSearch: (i) the predicted concept-bearing tokens and (ii) the predicted concept-to-code alignments.

- **Highlight Accuracy.** We use $\delta_{highlight}$ to threshold whether a token is predicted as concept-bearing. We then measure the **Precision** and **Recall** of these predicted tokens, using our annotated labels as the ground truth. The final evaluation metric is the average precision and recall across CodeSearchNet testing split.
- **Alignment Accuracy.** To evaluate alignment, we compute the cosine similarity between each query concept centroid (Equation 7) and all code line embeddings (Equation 8). If a code line’s similarity exceeds a predefined threshold δ_{align} , it is considered an aligned candidate. We then measure the average alignment **Precision** and **Recall** by comparing whether the predicted aligned code lines are within the annotated ground-truth alignments. Additionally, we report **Recall@k**, which evaluates whether the top-k most similar code lines retrieved for each query concept token include any of the annotated ground-truth alignments.

4.5.2 Results. As shown in Figure 8, XSearch achieves consistently high performance across all six programming languages in both alignment and highlight evaluation.

For concept-to-code alignment, the model achieves high top-3 recall (Alignment Recall@3 > 0.88), meaning that the correct code statements are usually ranked among the top candidates. For concept highlighting, both precision and recall are high for query and code tokens, indicating that the model reliably identifies tokens that correspond to meaningful concepts, in line with human annotations.

Taken together, these results indicate that the explanations produced by XSearch are faithful to the ground truth: both highlighted tokens and aligned code statements closely match the ground-truths. This level of faithfulness supports the use of concept alignment as a reliable explanation mechanism and provides basis for our user study in Section 4.7.

Table 5. Ablation Study of *XSearch*. Best scores are in **bold**.

Components				Languages (MRR)					
Alignment Loss (In-sample Neg.)	Alignment Loss (Cross-sample Neg.)	Highlight Loss	Type Emb.	Ruby	JS	Go	Python	Java	PHP
✓	✓	✗	✗	0.647	0.586	0.860	0.687	0.672	0.511
✗	✗	✓	✗	0.002	0.008	0.003	0.004	0.003	0.001
✗	✓	✓	✗	0.636	0.578	0.847	0.689	0.646	0.513
✓	✗	✓	✗	0.017	0.034	0.024	0.012	0.090	0.022
✓	✓	✗	✓	0.668	0.577	0.857	0.701	0.667	0.551
✓	✓	✓	✓	0.725	0.654	0.892	0.720	0.690	0.604

4.6 RQ3: Ablation Study

4.6.1 **Setup.** We conduct ablation studies to assess the contributions of alignment loss and highlight loss. We consider the following variants:

- **Alignment loss only:** The model is trained with only the alignment loss without highlight loss, i.e. all tokens are assumed to be highlighted.
- **Highlight loss only:** The model is trained with only the highlight loss, refining token importance and highlight distribution.
- **Alignment loss w/o negative sampling + Highlight loss:** The model is trained on both losses, but all negatives participate in alignment loss computation.
- **Alignment loss w/o cross-sample negatives + Highlight loss:** The model is trained on both losses, but all negatives come from the intra-sample code (Section 3.3.2).
- **Alignment loss + Highlight loss w/o type embedding:** The model is trained on both losses, but without adding data type embeddings to the code tokens' hidden states.
- **Full setting:** The model is trained with both alignment loss and highlight loss, evaluating their combined effect on overall performance.

And we observe the MRR performance of different settings.

4.6.2 **Results.** Table 5 summarizes the ablation results across six programming languages. The full setting, which combines alignment loss, highlight loss, cross-sample negative sampling, and type embeddings, consistently achieves the best performance. Removing any component leads to a clear and consistent drop in retrieval accuracy.

Using alignment loss alone degrades performance, suggesting that alignment benefits from explicit supervision on which tokens correspond to meaningful concepts. In contrast, relying only on highlight loss results in an even larger decline, indicating that identifying important tokens without enforcing their correspondence to code is insufficient for effective retrieval.

We also observe a substantial performance drop when cross-sample negative sampling is removed from the alignment loss. This suggests that hard negatives drawn across samples play an important role in preventing mis-alignments. Overall, these results show that all the components of *XSearch* work together in a complementary manner.

4.7 RQ4: User Study on Explanation

We conduct a controlled user study to test whether *XSearch*'s concept-alignment explanations help developers make (i) **faster** and (ii) more **accurate** accept or reject decisions when choosing among plausible code candidates.

4.7.1 **Participants.** We recruited 20 participants from four top universities with different backgrounds, including undergraduate students, master's students, PhD candidates, and research staff. Before the study, each participant was asked to complete a short questionnaire covering their (i)

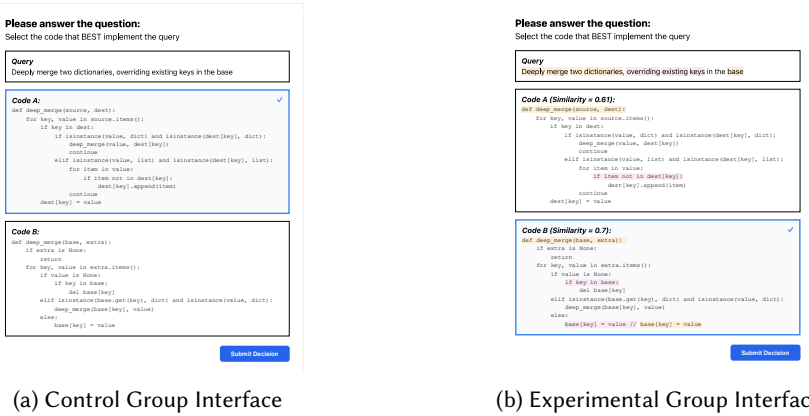


Fig. 9. The query asks for an in-place merge of two dictionaries and overwrites existing keys in the base dictionary. Code A didn’t overwrite existing keys, so the correct answer is Code B. Without explanations, users might omit such details and select the wrong answer.

academic level, (ii) years of programming, and (iii) self-reported Python familiarity. We then formed 10 matched pairs with similar profiles and assigned one participant to the **experimental group (EG)** and the other to the **control group (CG)**. As a result, two groups are of equal size ($n = 10$), with the same demographic distribution.

4.7.2 Task Design. We constructed 20 evaluation tasks from CodeSearchNet in **Python**. Each task paired a natural-language query with **five** candidates, only one of which matched the query intent. Tasks spanned five functional categories (Table 6) and three difficulty levels based on Line-of-Code (LOC) and branch complexity: Easy (< 8 LOC), Medium (8–16 LOC), and Difficult (> 16 LOC).

Table 6. Task Distribution by Functional Category and Examples

Category	Count	Examples
System and Configuration Management	5	<i>Removes an environment variable</i>
Data Conversion	5	<i>Convert any value to a boolean</i>
Data Parsing	4	<i>Extract the parts of a date given a timestamp</i>
Data I/O	4	<i>Writes the experimental setup to a JSON file</i>
Data Validation	2	<i>Check allowed file extensions</i>
Total	20	–

4.7.3 UI Interface. The CG interface (Figure 9a) mimics standard search, displaying only candidate code. The EG interface (Figure 9b) additionally presents XSearch’s explanations: (i) highlighted query concepts and (ii) their corresponding aligned code spans in each candidate. Candidate ordering and ground-truth answers remained identical across groups.

4.7.4 Study Setup. Before starting the main session of the user study, all participants complete a warm-up session consisting of 5 practice tasks. This session allows participants to familiarize themselves with the task format and interface without contributing to the evaluation results. Both groups receive the same warm-up tasks to ensure equal preparation. In the main session, we evaluate the participants’ performance along two perspectives:

- **Accuracy**, i.e., whether the participant selects the correct code snippet that matches the query’s intended functionality;

- **Completion Time**, i.e., the time taken to make a selection for the main session.

To compare the performance between the CG and EG, we use the Mann-Whitney U test [54]. This test is appropriate for independent samples when the assumption of normality may not hold, and it evaluates whether there is a significant difference in the distributions of the two groups.

4.7.5 **Results.** The results are summarized in Table 7, with the following observation:

- **Accuracy:** The experimental group (EG), achieves an average accuracy of **80.00%**, compared to **69.50%** in the control group (CG), showing an improvement of **10.5%**. The Mann-Whitney U test reveals that this improvement is statistically significant ($p = 0.0297 < 0.05$), indicating that *XSearch* substantially improves accuracy.
- **Efficiency:** The EG group completes the tasks faster, with an average task time of **1775.2** seconds compared to **2863.4** seconds in the CG group, resulting in a **38%** reduction in time. The Mann-Whitney U test also shows this difference is statistically significant ($p = 0.0188 < 0.05$), indicating a clear efficiency gain from the enhanced interface.

Table 7. Performance Comparison: Novice vs. Expert Groups

User Experience Level	Accuracy (%)		Time (s)		Improvement	
	CG	EG	CG	EG	Δ Acc	Time Saved
Low (Novice)	75.00	76.25	2811.7	2162.3	+1.25	23.1%
Medium (Intermediate)	45.00	85.00	3336.0	1300.0	+40.00	61.0%
High (Expert)	66.67	82.00	2809.3	1560.6	+15.33	44.5%
Average	69.50	80.00	2863.4	1775.2	+10.5%	38.0%

How highlighting helps users? Our results confirm that concept highlighting significantly reduces cognitive load across all users' programming experience levels. Notably, we observed a **leveling up** effect among **Intermediate** users. With *XSearch*'s explanations, they achieved higher accuracy than Expert users in the control group (85% vs. 66%) while requiring less than half the time (1300s vs. 2800s). This suggests the tool effectively bridges the expertise gap by directing attention to critical logic instead of requiring line-by-line reading.

Why does EG achieve higher accuracy? Post-study interviews with CG (Control Group) participants revealed that incorrect choices were primarily due to **keyword match** bias. Participants often chose an incorrect candidate simply because it contained some keywords similar to the query, despite missing the core logic. As shown in the case in Figure 9, the correct and incorrect snippets share significant keyword overlaps. Without the concept-to-line alignments provided by *XSearch*, CG participants failed to identify where the logic differs from the query's intent (e.g., Only Code B implements "overriding" existing keys).

5 Threats to Validity

Internal threats. Concept alignment labels are generated using GPT-4o [61] and may contain semantic noise. We mitigate this risk through formatting constraints and manual verification, though residual noise may remain. Our goal is not perfect annotation, but to have richer supervision through label augmentation. Model performance may also depend on hyperparameter choices. However, our sensitivity analysis (Figure 4) shows stable performance within reasonable ranges.

External threats. Our evaluation is based on CodeSearchNet and CoSQA+, which may not cover all real-world code search scenarios, such as large repository-level retrieval. Nevertheless, these benchmarks span six languages and are widely used in prior works, enabling fair comparison. Our user study involves a limited number of participants under controlled settings, further real-world studies are needed to assess long-term usability.

6 Conclusion

In this work, we present *XSearch*, an intrinsically explainable code retrieval model that addresses two key limitations of existing approaches: limited explainability and poor generalization under distribution shift. Unlike traditional methods that rely on global embedding search, *XSearch* explicitly formulates retrieval as a concept-alignment task. This reframes retrieval from an inductive to a deductive perspective: the model is designed to (i) identify the functional requirements in a query and (ii) verify how many of those requirements are satisfied by a candidate code snippet. Extensive experiments show that *XSearch* achieves competitive performance on in-distribution benchmarks and outperforms state-of-the-art models by up to 15× on out-of-distribution datasets. In addition, our user study demonstrates that concept-alignment explanations improve decision speed by 38% and accuracy by 10.5% when users accept or reject code recommendations. Overall, *XSearch* bridges the gap between performance and interpretability in code search, offering a scalable, user-centric solution for semantic retrieval.

Data Availability

The datasets and code used in this study are available at <https://anonymous.4open.science/r/Xsearch-2EC0>.

References

- [1] Wasi Uddin Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. 2020. A transformer-based approach for source code summarization. *arXiv preprint arXiv:2005.00653* (2020).
- [2] Anonymous. 2025. XSearch. <https://sites.google.com/view/xai-search/home> Accessed: 2025-03-15.
- [3] Suborno Deb Bappon, Saikat Mondal, and Banani Roy. 2024. AUTOGENICS: Automated Generation of Context-Aware Inline Comments for Code Snippets on Programming Q&A Sites Using LLM. In *2024 IEEE International Conference on Source Code Analysis and Manipulation (SCAM)*. IEEE, 24–35.
- [4] Kaj Bostrom, Harsh Jhamtani, Hao Fang, Sam Thomson, Richard Shin, Patrick Xia, Benjamin Van Durme, Jason Eisner, and Jacob Andreas. 2024. Language-to-Code Translation with a Single Labeled Example. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. 8101–8112.
- [5] Wing-Kwan Chan, Hong Cheng, and David Lo. 2012. Searching connected API subgraph via text phrases. In *Proceedings of the ACM SIGSOFT 20th international symposium on the foundations of software engineering*. 1–11.
- [6] Gong Chen, Xiaoyuan Xie, Daniel Tang, Qi Xin, and Wenjie Liu. 2024. HedgeCode: A Multi-Task Hedging Contrastive Learning Framework for Code Search. In *2025 IEEE/ACM 47th International Conference on Software Engineering (ICSE)*. IEEE Computer Society, 89–100.
- [7] Junkai Chen, Xing Hu, Zhenhao Li, Cuiyun Gao, Xin Xia, and David Lo. 2024. Code search is all you need? improving code suggestions with code search. In *Proceedings of the IEEE/ACM 46th international conference on software engineering*. 1–13.
- [8] Mouxiang Chen, Hao Tian, Zhongxin Liu, Xiaoxue Ren, and Jianling Sun. 2024. Jumpcoder: Go beyond autoregressive code via online modification. *arXiv preprint arXiv:2401.07870* (2024).
- [9] Xinyun Chen, Chang Liu, and Dawn Song. 2018. Tree-to-tree neural networks for program translation. *Advances in neural information processing systems* 31 (2018).
- [10] Yuxuan Chen, Guangsheng Ou, Mingwei Liu, Yanlin Wang, and Zibin Zheng. 2024. Are Decoder-Only Large Language Models the Silver Bullet for Code Search? *arXiv preprint arXiv:2410.22240* (2024).
- [11] Zhenlong Dai, Chang Yao, WenKang Han, Ying Yuan, Zhipeng Gao, and Jingyuan Chen. 2024. Mpcoder: Multi-user personalized code generator with explicit and implicit style representation learning. *arXiv preprint arXiv:2406.17255* (2024).
- [12] Luca Di Grazia and Michael Pradel. 2023. Code search: A survey of techniques for finding code. *Comput. Surveys* 55, 11 (2023), 1–31.
- [13] Aleksandra Eliseeva, Yaroslav Sokolov, Egor Bogomolov, Yaroslav Golubev, Danny Dig, and Timofey Bryksin. 2023. From commit message generation to history-aware commit message completion. In *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 723–735.
- [14] Guodong Fan, Shizhan Chen, Cuiyun Gao, Jianmao Xiao, Tao Zhang, and Zhiyong Feng. 2024. Rapid: Zero-shot domain adaptation for code search with pre-trained models. *ACM Transactions on Software Engineering and Methodology* 33, 5 (2024), 1–35.
- [15] Zhiyu Fan, Xiang Gao, Martin Mirchev, Abhik Roychoudhury, and Shin Hwei Tan. 2023. Automated repair of programs from large language models. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 1469–1481.
- [16] Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, et al. 2020. Codebert: A pre-trained model for programming and natural languages. *arXiv preprint arXiv:2002.08155* (2020).
- [17] Michael Fu, Chakkrit Tantithamthavorn, Trung Le, Van Nguyen, and Dinh Phung. 2022. VulRepair: a T5-based automated software vulnerability repair. In *Proceedings of the 30th ACM joint european software engineering conference and symposium on the foundations of software engineering*. 935–947.
- [18] Jing Gong, Yanghui Wu, Linxi Liang, Zibin Zheng, and Yanlin Wang. 2024. CoSQA+: Enhancing Code Search Dataset with Matching Code. *arXiv preprint arXiv:2406.11589* (2024).
- [19] Wenchao Gu, Zongyi Lyu, Yanlin Wang, Hongyu Zhang, Cuiyun Gao, and Michael R. Lyu. 2025. SPENCER: Self-Adaptive Model Distillation for Efficient Code Retrieval. *ACM Transactions on Software Engineering and Methodology* (2025).
- [20] Wenchao Gu, Yanlin Wang, Lun Du, Hongyu Zhang, Shi Han, Dongmei Zhang, and Michael Lyu. 2022. Accelerating code search with deep hashing and code classification. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 2534–2544.

- [21] Xiaodong Gu, Hongyu Zhang, and Sunghun Kim. 2018. Deep code search. In *Proceedings of the 40th international conference on software engineering*. 933–944.
- [22] Daya Guo, Shuai Lu, Nan Duan, Yanlin Wang, Ming Zhou, and Jian Yin. 2022. Unixcoder: Unified cross-modal pre-training for code representation. *arXiv preprint arXiv:2203.03850* (2022).
- [23] Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, Shujie Liu, Long Zhou, Nan Duan, Alexey Svyatkovskiy, Shengyu Fu, et al. 2020. Graphcodebert: Pre-training code representations with data flow. *arXiv preprint arXiv:2009.08366* (2020).
- [24] Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Yu Wu, YK Li, et al. 2024. DeepSeek-Coder: When the Large Language Model Meets Programming—The Rise of Code Intelligence. *arXiv preprint arXiv:2401.14196* (2024).
- [25] Vincent J Hellendoorn, Charles Sutton, Rishabh Singh, Petros Maniatis, and David Bieber. 2019. Global relational models of source code. In *International conference on learning representations*.
- [26] Dan Hendrycks, Xiaoyuan Liu, Eric Wallace, Adam Dziedzić, Rishabh Krishnan, and Dawn Song. 2020. Pretrained transformers improve out-of-distribution robustness. *arXiv preprint arXiv:2004.06100* (2020).
- [27] Emily Hill, Manuel Roldan-Vega, Jerry Alan Fails, and Greg Mallet. 2014. NL-based query refinement and contextualized code search results: A user study. In *2014 Software Evolution Week-IEEE Conference on Software Maintenance, Reengineering, and Reverse Engineering (CSMR-WCRE)*. IEEE, 34–43.
- [28] Baizhou Huang, Shuai Lu, Weizhu Chen, Xiaojun Wan, and Nan Duan. 2023. Enhancing large language models in coding through multi-perspective self-consistency. *arXiv preprint arXiv:2309.17272* (2023).
- [29] Yufan Huang, Mengnan Qi, Yongqiang Yao, Maoquan Wang, Bin Gu, Colin Clement, and Neel Sundaresan. 2023. Program translation via code distillation. *arXiv preprint arXiv:2310.11476* (2023).
- [30] Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, et al. 2024. Qwen2.5-coder technical report. *arXiv preprint arXiv:2409.12186* (2024).
- [31] Faria Huq, Masum Hasan, Md Mahim Anjum Haque, Sazan Mahbub, Anindya Iqbal, and Toufique Ahmed. 2022. Review4repair: Code review aided automatic program repairing. *Information and Software Technology* 143 (2022), 106765.
- [32] Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. 2019. Codesearchnet challenge: Evaluating the state of semantic code search. *arXiv preprint arXiv:1909.09436* (2019).
- [33] Jeevana Priya Inala, Chenglong Wang, Mei Yang, Andres Cudas, Mark Encarnación, Shuvendu Lahiri, Madanlal Musuvathi, and Jianfeng Gao. 2022. Fault-aware neural code rankers. *Advances in Neural Information Processing Systems* 35 (2022), 13419–13432.
- [34] Chen Ji, Su Yang, Hongyu Sun, and Yuqing Zhang. 2024. Applying Contrastive Learning to Code Vulnerability Type Classification. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. 11942–11952.
- [35] Siyuan Jiang, Jia Li, He Zong, Huanyu Liu, Hao Zhu, Shukai Hu, Erlu Li, Jiazheng Ding, Yu Han, Wei Ning, et al. 2024. aixcoder-7b: A lightweight and effective large language model for code completion. *arXiv e-prints* (2024), arXiv–2410.
- [36] Zhonghao Jiang, Xiaoxue Ren, Meng Yan, Wei Jiang, Yong Li, and Zhongxin Liu. 2025. CoSIL: Software Issue Localization via LLM-Driven Code Repository Graph Searching. *arXiv preprint arXiv:2503.22424* (2025).
- [37] Tae-Hwan Jung. 2021. Commitbert: Commit message generation using pre-trained programming language model. *arXiv preprint arXiv:2105.14242* (2021).
- [38] Sungmin Kang, Louis Milliken, and Shin Yoo. 2024. Identifying inaccurate descriptions in llm-generated code comments via test execution. *arXiv preprint arXiv:2406.14836* (2024).
- [39] Kisub Kim, Dongsun Kim, Tegawendé F Bissyandé, Eunjong Choi, Li Li, Jacques Klein, and Yves Le Traon. 2018. FaCoY: a code-to-code search engine. In *Proceedings of the 40th International Conference on Software Engineering*. 946–957.
- [40] Marie-Anne Lachaux, Baptiste Roziere, Lowik Chanussot, and Guillaume Lample. 2020. Unsupervised translation of programming languages. *arXiv preprint arXiv:2006.03511* (2020).
- [41] Haochen Li, Chunyan Miao, Cyril Leung, Yanxin Huang, Yuan Huang, Hongyu Zhang, and Yanlin Wang. 2022. Exploring representation-level augmentation for code search. *arXiv preprint arXiv:2210.12285* (2022).
- [42] Jiawei Li, David Faragó, Christian Petrov, and Iftekhar Ahmed. 2025. Optimization is Better than Generation: Optimizing Commit Message Leveraging Human-written Commit Message. *arXiv preprint arXiv:2501.09861* (2025).
- [43] Lingwei Li, Li Yang, Huaxi Jiang, Jun Yan, Tiejian Luo, Zihan Hua, Geng Liang, and Chun Zuo. 2022. AUGER: automatically generating review comments with pre-training models. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1009–1021.
- [44] Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, et al. 2023. Starcoder: may the source be with you! *arXiv preprint arXiv:2305.06161* (2023).
- [45] Wen-Ding Li and Kevin Ellis. 2025. Is programming by example solved by llms? *Advances in Neural Information Processing Systems* 37 (2025), 44761–44790.

- [46] Xiaonan Li, Yeyun Gong, Yelong Shen, Xipeng Qiu, Hang Zhang, Bolun Yao, Weizhen Qi, Daxin Jiang, Weizhu Chen, and Nan Duan. 2022. Coderetriever: A large scale contrastive pre-training method for code search. In *Proceedings of the 2022 conference on empirical methods in natural language processing*. 2898–2910.
- [47] Zhen Li, Deqing Zou, Shouhuai Xu, Hai Jin, Yawei Zhu, and Zhaoxuan Chen. 2021. Sysevr: A framework for using deep learning to detect software vulnerabilities. *IEEE Transactions on Dependable and Secure Computing* 19, 4 (2021), 2244–2258.
- [48] Zhen Li, Deqing Zou, Shouhuai Xu, Xinyu Ou, Hai Jin, Sujuan Wang, Zhijun Deng, and Yuyi Zhong. 2018. Vuldeepecker: A deep learning-based system for vulnerability detection. *arXiv preprint arXiv:1801.01681* (2018).
- [49] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*. 2980–2988.
- [50] Chao Liu, Xin Xia, David Lo, Cuiyun Gao, Xiaohu Yang, and John Grundy. 2021. Opportunities and challenges in code search tools. *ACM Computing Surveys (CSUR)* 54, 9 (2021), 1–40.
- [51] Wei Liu, Ailun Yu, Daoguang Zan, Bo Shen, Wei Zhang, Haiyan Zhao, Zhi Jin, and Qianxiang Wang. 2024. Graphcoder: Enhancing repository-level code completion via code context graph-based retrieval and language model. *arXiv preprint arXiv:2406.07003* (2024).
- [52] Meili Lu, Xiaobing Sun, Shaowei Wang, David Lo, and Yucong Duan. 2015. Query expansion via wordnet for effective code search. In *2015 IEEE 22nd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*. IEEE, 545–549.
- [53] Fei Lv, Hongyu Zhang, Jian-guang Lou, Shaowei Wang, Dongmei Zhang, and Jianjun Zhao. 2015. Codehow: Effective code search based on api understanding and extended boolean model (e). In *2015 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 260–270.
- [54] Henry B Mann and Donald R Whitney. 1947. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics* (1947), 50–60.
- [55] Collin McMillan, Mark Grechanik, Denys Poshyvanyk, Qing Xie, and Chen Fu. 2011. Portfolio: finding relevant functions and their usage. In *Proceedings of the 33rd International Conference on Software Engineering*. 111–120.
- [56] Microsoft. 2021. GraphCodeBERT model. <https://huggingface.co/microsoft/graphcodebert-base>.
- [57] Ansong Ni, Srini Iyer, Dragomir Radev, Veselin Stoyanov, Wen-tau Yih, Sida Wang, and Xi Victoria Lin. 2023. Lever: Learning to verify language-to-code generation with execution. In *International Conference on Machine Learning*. PMLR, 26106–26128.
- [58] Liming Nie, He Jiang, Zhilei Ren, Zeyi Sun, and Xiaochen Li. 2016. Query expansion based on crowd knowledge for code search. *IEEE Transactions on Services Computing* 9, 5 (2016), 771–783.
- [59] Yu Nong, Rainy Sharma, Abdelwahab Hamou-Lhadj, Xiapu Luo, and Haipeng Cai. 2022. Open science in software engineering: A study on deep learning-based vulnerability detection. *IEEE Transactions on Software Engineering* 49, 4 (2022), 1983–2005.
- [60] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018).
- [61] OpenAI. 2024. GPT-4o Technical Report. <https://openai.com/index/hello-gpt-4o/>. Accessed: 2025-01.
- [62] Yun Peng, Shuzheng Gao, Cuiyun Gao, Yintong Huo, and Michael Lyu. 2024. Domain knowledge matters: Improving prompts with fix templates for repairing python type errors. In *Proceedings of the 46th IEEE/ACM international conference on software engineering*. 1–13.
- [63] David Piorkowski, Austin Z Henley, Tahmid Nabi, Scott D Fleming, Christopher Scaffidi, and Margaret Burnett. 2016. Foraging and navigations, fundamentally: developers’ predictions of value and cost. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. 97–108.
- [64] Mukund Raghthaman, Yi Wei, and Youssef Hamadi. 2016. SWIM: synthesizing what I mean: code search and idiomatic snippet synthesis. In *Proceedings of the 38th International Conference on Software Engineering*. 357–367.
- [65] Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950* (2023).
- [66] Baptiste Roziere, Jie M Zhang, Francois Charton, Mark Harman, Gabriel Synnaeve, and Guillaume Lample. 2021. Leveraging automated unit tests for unsupervised code translation. *arXiv preprint arXiv:2110.06773* (2021).
- [67] Anthony Saieva, Saikat Chakraborty, and Gail Kaiser. 2024. Reinforeset: Reinforcing semantic code similarity for cross-lingual code search models. In *2024 IEEE International Conference on Source Code Analysis and Manipulation (SCAM)*. IEEE, 177–188.
- [68] Chaochen Shi, Borui Cai, Yao Zhao, Longxiang Gao, Keshav Sood, and Yong Xiang. 2023. Coss: Leveraging statement semantics for code summarization. *IEEE Transactions on Software Engineering* 49, 6 (2023), 3472–3486.
- [69] Ensheng Shi, Yanlin Wang, Wenchao Gu, Lun Du, Hongyu Zhang, Shi Han, Dongmei Zhang, and Hongbin Sun. 2023. Cocosoda: Effective contrastive learning for code search. In *2023 IEEE/ACM 45th International Conference on Software*

- Engineering (ICSE)*. IEEE, 2198–2210.
- [70] Jaspreet Singh and Avishek Anand. 2019. Exs: Explainable search using local model agnostic interpretability. In *Proceedings of the twelfth ACM international conference on web search and data mining*. 770–773.
- [71] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2020. MPNet: Masked and Permuted Pre-training for Language Understanding. *arXiv preprint arXiv:2004.09297* (2020).
- [72] Benjamin Steenhoek, Md Mahbubur Rahman, Richard Jiles, and Wei Le. 2023. An empirical study of deep learning models for vulnerability detection. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2237–2248.
- [73] Elias Stengel-Eskin, Archiki Prasad, and Mohit Bansal. 2024. Regal: Refactoring programs to discover generalizable abstractions. *arXiv preprint arXiv:2401.16467* (2024).
- [74] Chia-Yi Su and Collin McMillan. 2024. Distilled GPT for source code summarization. *Automated Software Engineering* 31, 1 (2024), 22.
- [75] Marc Szafraniec, Baptiste Roziere, Hugh Leather, Francois Charton, Patrick Labatut, and Gabriel Synnaeve. 2022. Code translation with compiler representations. *arXiv preprint arXiv:2207.03578* (2022).
- [76] Xunzhu Tang, Saad Ezzini, Haoye Tian, Yewei Song, Jacques Klein, Tegawende F Bissyande, et al. 2023. Hyperbolic code retrieval: a novel approach for efficient code search using hyperbolic space embeddings. *arXiv preprint arXiv:2308.15234* (2023).
- [77] Xunzhu Tang, Kisub Kim, Yewei Song, Cedric Lothritz, Bei Li, Saad Ezzini, Haoye Tian, Jacques Klein, and Tegawendé F Bissyandé. 2024. CodeAgent: Autonomous Communicative Agents for Code Review. *arXiv preprint arXiv:2402.02172* (2024).
- [78] Ze Tang, Xiaoyu Shen, Chuanyi Li, Jidong Ge, Liguang Huang, Zhelin Zhu, and Bin Luo. 2022. AST-trans: Code summarization with efficient tree-structured attention. In *Proceedings of the 44th International Conference on Software Engineering*. 150–162.
- [79] Ali Tehranijamsaz, Arijit Bhattacharjee, Le Chen, Nesreen K Ahmed, Amir Yazdanbakhsh, and Ali Jannesari. 2024. CodeRosetta: Pushing the Boundaries of Unsupervised Code Translation for Parallel Programming. *arXiv preprint arXiv:2410.20527* (2024).
- [80] Rosalia Tufano, Simone Masiero, Antonio Mastropaolo, Luca Pascarella, Denys Poshyvanyk, and Gabriele Bavota. 2022. Using pre-trained models to boost code review automation. In *Proceedings of the 44th international conference on software engineering*. 2291–2302.
- [81] Chong Wang, Xin Peng, Zhenchang Xing, Yue Zhang, Mingwei Liu, Rong Luo, and Xiujie Meng. 2023. Xcos: Explainable code search based on query scoping and knowledge graph. *ACM Transactions on Software Engineering and Methodology* 32, 6 (2023), 1–28.
- [82] Chengpeng Wang, Wuqi Zhang, Zian Su, Xiangzhe Xu, Xiaoheng Xie, and Xiangyu Zhang. 2024. LLMDFa: analyzing dataflow in code with large language models. *Advances in Neural Information Processing Systems* 37 (2024), 131545–131574.
- [83] Xin Wang, Yasheng Wang, Fei Mi, Pingyi Zhou, Yao Wan, Xiao Liu, Li Li, Hao Wu, Jin Liu, and Xin Jiang. 2021. Syncobert: Syntax-guided multi-modal contrastive pre-training for code representation. *arXiv preprint arXiv:2108.04556* (2021).
- [84] Yue Wang, Hung Le, Akhilesh Deepak Gotmare, Nghi DQ Bui, Junnan Li, and Steven CH Hoi. 2023. Codet5+: Open code large language models for code understanding and generation. *arXiv preprint arXiv:2305.07922* (2023).
- [85] Yue Wang, Weishi Wang, Shafiq Joty, and Steven CH Hoi. 2021. Codet5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation. *arXiv preprint arXiv:2109.00859* (2021).
- [86] Xin Xia, Lingfeng Bao, David Lo, Zhenchang Xing, Ahmed E Hassan, and Shanping Li. 2017. Measuring program comprehension: A large-scale field study with professionals. *IEEE Transactions on Software Engineering* 44, 10 (2017), 951–976.
- [87] Pengyu Xue, Linhao Wu, Zhongxing Yu, Zhi Jin, Zhen Yang, Xinyi Li, Zhenyu Yang, and Yue Tan. 2024. Automated commit message generation with large language models: An empirical study and beyond. *IEEE Transactions on Software Engineering* (2024).
- [88] Michihiro Yasunaga and Percy Liang. 2021. Break-it-fix-it: Unsupervised learning for program repair. In *International conference on machine learning*. PMLR, 11941–11952.
- [89] Xinran Yu, Chun Li, Minxue Pan, and Xuandong Li. 2024. DroidCoder: Enhanced Android Code Completion with Context-Enriched Retrieval-Augmented Generation. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*. 681–693.
- [90] Fengji Zhang, Bei Chen, Yue Zhang, Jacky Keung, Jin Liu, Daoguang Zan, Yi Mao, Jian-Guang Lou, and Weizhu Chen. 2023. Repocoder: Repository-level code completion through iterative retrieval and generation. *arXiv preprint arXiv:2303.12570* (2023).

- [91] Guodong Zhang, Tianyu Yao, Jiawei Qin, Yitao Li, Qiao Ma, and Donghong Sun. 2025. CodeSAGE: A multi-feature fusion vulnerability detection approach using code attribute graphs and attention mechanisms. *Journal of Information Security and Applications* 89 (2025), 103973.
- [92] Qunjun Zhang, Chunrong Fang, YE Shang, Tongke Zhang, Shengcheng Yu, and Zhenyu Chen. 2024. No man is an island: Towards fully automatic programming by code search, code generation and program repair. *arXiv preprint arXiv:2409.03267* (2024).