

Efficient and Scalable Provenance Tracking for LLM-Generated Code Snippets

Andrea Gurioli^a, Davide D’Ascenzo^b, Federico Pennino^a, Maurizio Gabbrielli^a, Stefano Zacchioli^c

^a*Dipartimento di Informatica – Scienza e Ingegneria, Università di Bologna, Bologna, Italy*

^b*Department of Control and Computer Engineering, Politecnico di Torino, Turin, Italy*

^c*Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France*

Abstract

Large language models (LLMs) for code completion and generation are increasingly used in software development, yet they may reproduce training examples verbatim and without authorship attribution, raising legal and ethical concerns around plagiarism and license compliance. Classical fingerprint-based plagiarism detectors, such as Winnowing, remain highly effective, yet the inspection requires comparing fragments of code to the entire training set, and their linear-time search makes them impractical for the billion-scale corpora used to train modern code LLMs. To bridge this gap, we introduce SOURCETRACKER, a 300M-parameter encoder tailored for code retrieval, together with a hybrid two-stage provenance-tracking pipeline HYBRIDSOURCETRACKER (HST). HST first narrows down a small set of candidate snippets via vector search, then re-ranks those candidates using Winnowing on exact fingerprints. We train and evaluate our system on a 10M-snippet subset of the THESTACKV2 dataset, with both verbatim and adapted snippets that emulate realistic identifier renaming. On an *in vitro* 100k-snippet search space with adapted queries, our hybrid approach reaches a mean reciprocal rank on par with Winnowing for 30-token fragments. Then, starting from windows ≥ 60 tokens, it consistently over-performs by up to

Email addresses: andrea.gurioli5@unibo.it (Andrea Gurioli),
davide.dascenzo@polito.it (Davide D’Ascenzo), federico.pennino2@unibo.it
(Federico Pennino), maurizio.gabbrielli@unibo.it (Maurizio Gabbrielli),
stefano.zacchioli@telecom-paris.fr (Stefano Zacchioli)

5.4% while preserving logarithmic-time query complexity. In a complementary evaluation using an LLM-based judge, we find that many retrieved snippets not labeled as ground truth are still highly similar to the expected sources, particularly with longer context windows, and thus remain useful for end users. Overall, our results demonstrate that integrating vector search with fingerprinting enables scalable, high-precision provenance tracking for code produced by LLMs, provided that their training data is made accessible. More broadly, our findings underscore the growing need for transparency around LLM training corpora as a prerequisite for effective provenance tracking, attribution, and license compliance.

Keywords: Provenance tracking, vector database, code generation, fingerprinting algorithm, code similarity, authorship attribution

1. Introduction




Figure 1: Overview of the HST architecture. We first encode the code fragment using SOURCETRACKER, extract the top 100 results from Qdrant, and then apply the Winnowing algorithm to rerank these results and identify the most similar code snippets.

The rise of Large Language Models (LLMs) for code completion and generation—such as Codex, CodeLlama, and StarCoder [1, 2, 3]—has become increasingly relevant as a tool for day-to-day code development [4]. Several authors [1, 2, 3] have highlighted that the performance of these models is notably driven by the data used during training. Carlini et al. [5] have demonstrated that LLMS can memorize part of the training data. The memorization can result in paraphrasing, reiterating ideas, or reproducing verbatim




Figure 2: Training procedure of SOURCETRACKER. Starting from the original latent space, SOURCETRACKER is trained using the CLIP loss to bring related code fragments and their complete snippets closer together while pushing apart different snippets.

sections from the training dataset [6]—a phenomenon referred to as “recitation” [7]. Consequently, the generation process of LLMs raises significant concerns. Generating code without acknowledging the original authors not only obscures the provenance of code, but also constitutes an unfair practice that disregards their intellectual contributions [8]. Depending on the context, this might constitute plagiarism [9] as well as expose users of generated code to the risk of inadvertently violating software licenses, potentially resulting in copyright infringement.

These challenges underscore the need to ensure transparency, ethical attribution, and license compliance in LLM-based code generation. In light of these issues, proper provenance tracking of generated code has become increasingly critical to meeting both legal and ethical obligations.

Fingerprinting techniques, such as Winnowing [10], have been adopted to assess whether a snippet of code can be classified as plagiarized in algorithms like MOSS [11]—widely used to inspect hundreds of thousands of HTML pages for potentially suspicious similarities. These algorithms typically run in linear time with respect to the dataset size.

LLMs, trained on extensive collections of code snippets, have consequently complicated the tracking of original authors by broadening the search space. Recent pre-training datasets exceed the scale of billions of samples [12, 13, 14], making algorithms with linear time complexity inadequate.

The evaluation of modern author detection systems must emphasize both efficiency and scalability. Traditional fingerprinting algorithms [11], while resilient to noise [10], must be adapted or complemented by new methods

capable of operating at the scale introduced by modern generative models.

Liu et al. [15] pioneered this task on a new scale of data by applying the infinity-gram retrieval algorithm from generated text to the training set, achieving a logarithmic time complexity within an average of 4.5 seconds per query. While these results improve speed, exact-match algorithms remain highly sensitive to noise. This sensitivity makes them inadequate for plagiarism detection [10] and contexts where code needs to be adapted to the context, such as code completion and generation [9].

Based on these findings, we aim to design a system that effectively retrieves the most similar code snippets from our training set for a code fragment of arbitrary length, generated by an LLM. This system must efficiently scale to search a space containing billions of files, providing end users with information about the authors of retrieved snippets and facilitating transparent tracking of the generated code.

We propose HYBRIDSOURCETRACKER (HST), a hybrid two-stage retrieval methodology that achieves logarithmic search time complexity while maintaining high detection accuracy by sequentially combining vector database retrieval with Winnowing-based re-ranking. In our approach, we first encode code snippets into dense vector representations using our novel fine-tuned model SOURCETRACKER, which are then indexed using a vector database. At inference time, a code fragment is encoded and used to query the database to efficiently retrieve the 100 most similar snippets based on vector similarity. The Winnowing algorithm [10] is subsequently applied to this candidate set, re-ranking the snippets using syntactic hash matching.

This approach outperforms traditional fingerprinting algorithms [10] in both efficacy and efficiency. Our methodology has been trained and tested to identify both exact and non-exact matches by employing a frequent word replacement algorithm, demonstrating its resilience to context adaptation.

We hope that this innovative approach will help in promoting transparency in code generation.

1.1. Contributions

Our main contributions are the following:

- We release **SOURCETRACKER**, a 300M-parameter model for code provenance retrieval, fine-tuned to retrieve complete code snippets from small fragments.

- We introduce and analyze **HYBRIDSOURCETRACKER (HST)** a two-step retrieval methodology that sequentially applies SOURCETRACKER and the Winnowing algorithm, offering theoretical guarantees of **logarithmic** search time complexity alongside practical scalability validated by extensive benchmarking, all while achieving performance comparable to traditional fingerprinting techniques.

2. Background

This section presents the main concepts and terminology used in the following. Section 2.1 introduces and distinguishes the key concepts used throughout the article, including authorship attribution, plagiarism detection, and related notions. Section 2.2 introduces the idea of code similarity. Section 2.3 outlines the methodology behind the MOSS and winnowing algorithms, which are state-of-the-art tools for plagiarism detection and will later serve as a baseline for comparison.

2.1. Terminology and scope

Throughout this work, we use the terms provenance tracking, plagiarism detection, memorization auditing, and authorship attribution, each referring to a distinct but closely related concept.

We use provenance tracking to denote the task of identifying one or more candidate source artifacts in a known corpus from which a generated code fragment may originate or be derived.

Memorization auditing instead concerns whether and to what extent an LLM has retained information from its training data; while memorization motivates our work, HST does not measure memorization or establish that a retrieved similarity results from memorization.

Plagiarism detection concerns whether reuse constitutes plagiarism, a determination that may depend on attribution, licensing, context, and other normative considerations beyond code similarity; HST can provide evidence useful for such an assessment but does not itself classify code as plagiarized.

Finally, authorship attribution concerns assigning authorship to an artifact. HST does not infer authors from coding style; rather, when provenance is established, metadata associated with the retrieved source can support provenance-based source attribution.

2.2. Code clones

Recent research highlights multiple ways through which Large Language Models can memorize and reproduce parts of their training data, such as through verbatim repetition, paraphrasing, or rephrasing underlying ideas [5]. We relied on Roy et al. [16] definition of *code clones* to represent, at various levels of granularity, how an LLM can reproduce part of the training data. Code clones refer to code fragments that are similar to each other through various forms of replication or modification, defined by the code clones taxonomy [17]:

- Type 1 (Exact Clones): Identical code fragments except for variations in whitespace, comments, and layout. These represent verbatim copies of code.
- Type 2 (Renamed Clones): Syntactically identical fragments where identifiers, literals, types, or function names have been changed. This type commonly occurs when developers adapt code to fit specific domain requirements.
- Type 3 (Near-Miss Clones): Code fragments with added, deleted, or modified statements, representing more substantial modifications.
- Type 4 (Semantic Clones): Fragments that perform the same function using different syntactic implementations.

In the context of LLM-generated code, we refer to Type 1 clones as direct verbatim recitations, while Type 2 clones represent domain adaptation, in which the model replicates the training data while modifying variable and function names to better align with the user’s context.

2.3. Winnowing and the MOSS engine

The Winnowing algorithm [10] is a widely adopted fingerprinting technique for software plagiarism detection. It forms the foundation of the MOSS system by selecting representative hashes called *fingerprints* that remain robust to minor syntactic or formatting changes in code. This method serves both as a baseline for our work and as the final refinement in our two-step procedure.

The algorithm begins with code *canonicalization*, wherein comments, whitespace, and other non-essential tokens are removed to produce a normalized representation. This canonicalized code is segmented into overlapping

k -grams (sequences of k characters), each of which is hashed, using the last 16 bits of the SHA-1 digest. A sliding window of size w is then applied over the hash sequence, and the minimum hash value in each window is retained as a fingerprint. This guarantees that any matching substring will produce at least one common fingerprint between two documents.

The MOSS retrieval engine [10] relies on the Winnowing algorithm to optimize time search. The engine initially applies the Winnowing algorithm to each document, producing a set of fingerprints stored in an inverted index (hash to document IDs). At query time, the algorithm generates fingerprints for the query document, drops the most frequent hashes, and inspects the index in rarest-first order, prioritizing hashes in the inverted index list that point to fewer candidate documents. The process is conducted under fixed budgets to build a compact candidate set, which is then re-ranked using exact Jaccard similarity on fingerprint sets to identify the top-K matches.

3. Methodology

Our primary aim is to provide users with insight into potential sources of inspiration for LLM generation. To achieve this, we focus our analysis on code fragments that exhibit sufficient similarity to be reliably linked to specific source snippets. We examine two types of code clones, Type 1 and Type 2 (see Section 2.2), both adapted from their original definitions to align with the context of LLM code generation. We restrict the analysis to these two types because provenance requires a sufficiently strong syntactic link to a single identifiable source: for near-miss and semantic clones (Types 3 and 4), a generated fragment may be a transformation of one source or a generalisation over many, and a unique ground-truth source cannot be defined (see Section 8.1.1).

For **Type 1 clones**, we refine the standard definition to focus exclusively on verbatim replication of complete or partial code snippets from the training set. For **Type 2 clones**, we adjust the definition to reflect domain-specific adaptation typical in LLM applications, where the model generates syntactically similar code while altering identifiers to fit the user’s context. To simulate this behavior, we apply a frequent word replacement algorithm that probabilistically substitutes recurring identifiers with random strings.

Our approach is centered on a two-stage pipeline, namely HST, that integrates the efficiency of vector database retrieval to address the retrieval phase

Algorithm 1 Replace frequent words in code fragment

```
1: function REPFREQWORDS(code_fragment)
2:   word_counts  $\leftarrow$  Count occurrences of words in code_fragment
3:   for all (word, count) in word_counts do
4:     if length(word) > 3 and count > 2 then
5:       if RANDOM  $\leq$  0.2 then  $\triangleright$  Uniform draw in  $[0, 1)$ 
6:         random_str  $\leftarrow$  GETRANDOMSTRING()
7:         Replace all occurrences of word in code_fragment with ran-
           dom_str
8:       end if
9:     end if
10:  end for
11:  return code_fragment
12: end function

13: function GETRANDOMSTRING()
14:   digits  $\leftarrow$  4 characters sampled uniformly at random from  $\{0, \dots, 9\}$ 
15:   letters  $\leftarrow$  4 characters sampled uniformly at random from  $\{a-z, A-Z\}$ 
16:   return random permutation of the concatenation digits || letters
17: end function
```

from large pre-training datasets, along with the efficacy of the Winnowing algorithm.

Initially, we fine-tuned an encoder with a metric learning objective to facilitate vectorial retrieval, as detailed in Section 3.2. We then compare the results of this approach with those obtained from the standard Winnowing algorithm (see Section 3.3). Finally, as illustrated in Figure 1, HST implementation, described in Section 3.4, integrates the two methodologies described above.

3.1. Dataset

We rely on a randomly sampled subset of 10 000 000 snippets of code from THESTACKV2 [12]—an open-source dataset used to train STARCORDER2—which enables realistic analysis both in terms of data scale and dataset composition to extract our ground truth. We extract fragments of code directly from complete snippets within the extracted subset; this preserves ground truth, since each fragment remains explicitly tied back to its original source.

During training and evaluation, we either use the fragments verbatim (Clone Type 1) or rephrase them (Clone Type 2) using our domain adaptation algorithm (Algorithm 1). Specifically, in Algorithm 1, for each word longer than three characters, appearing more than twice in a code fragment, we apply a 20% probability of replacing all its occurrences with a randomly generated string, provided the word length exceeds three characters. This transformation process, formalized in Algorithm 1, emulates the domain-adaptation phenomenon in which LLMs modify variable and function names while preserving the underlying code structure. We randomly sampled 10% of the extracted subset as the test set for the final performance evaluation phase.

3.2. SOURCETRACKER

To investigate how deep learning techniques can enhance code provenance retrieval, we fine-tuned an encoder to retrieve syntactically similar code snippets from small code fragments. By utilizing a vectorial encoder representation, we can leverage vector databases. This approach significantly accelerates the search process, reducing query time complexity to logarithmic [18] compared to the dataset size.

We fine-tune a bidirectional encoder to align the vector representations of code fragments with their corresponding complete code snippets (see Figure 2). Building on the ModularStarEncoder (MoSE) [19]—a 1-billion-parameter multi-exit encoder architecture designed for code retrieval—we utilize the first 9 layers, totaling 300M parameters, to compute a 1024-dimensional vector representation for each code document. We opted for this configuration due to the best trade-off between retrieval effectiveness, reported in [19], and the computational and memory constraints of our hardware setup. Following the threshold of interest suggested by Ziegler [7], we extract fixed-length fragments of 60 tokens from each code snippet. To align our model with domain adaptation scenarios, we use adapted code fragments (Type 2 Clones) 50% of the time during training, ensuring robustness to identifier modifications while preserving structural similarity. For the remaining 50% of the training, we utilize Type 1 clones.

The model is trained using the CLIP (Contrastive Language-Image Pre-training) loss function [20], a contrastive objective that maximizes the cosine similarity between matching code fragment-snippet pairs while minimizing similarity for non-matching pairs. Specifically, the loss encourages the encoder to pull related fragments and their full snippets closer together in the learned latent space, while pushing apart embeddings from different snippets.

We employ a batch size of 2048 code fragment-snippet pairs for 2000 steps, enabling the contrastive mechanism to leverage a large number of negative examples per batch, which improves the discriminative quality of the learned representations.

During training, we use the AdamW optimizer with a learning rate of $6e-4$ and a warmup phase lasting 500 steps.

3.2.1. Lookup phase optimization

To accelerate similarity search in our system, we employ **Qdrant** [21], an open-source, high-performance vector database optimized for approximate nearest neighbor (ANN) retrieval. Qdrant is based on the *Hierarchical Navigable Small World (HNSW)* [18] algorithm, which constructs a multi-layer graph to efficiently navigate high-dimensional vector spaces. This method achieves an average search complexity of $\mathcal{O}(\log N)$, where N represents the number of indexed vectors, significantly reducing search latency compared to the MOSS algorithm [11], which maintains linear complexity.

3.3. Winnowing

To establish a solid baseline for comparison, we implemented the Winnowing algorithm as detailed in Section 2.3, following the standard MOSS configuration. This implementation serves as our primary non-learning baseline and represents the state of the art in syntactic code plagiarism detection.

We configured the Winnowing algorithm with the following parameters: a k -gram size of $k = 5$, a sliding window size of $w = 4$, and hashing each k -gram with SHA-1 while retaining the last 16 bits. We also set a retrieval budget of 64 hash inspections per query. For the high-frequency hash filtering step, we empirically determined a threshold to exclude hashes that appear in more than 0.5% of documents based on preliminary experiments. This threshold strikes a balance between reducing noise and maintaining enough fingerprints for effective matching. All candidates retrieved within the hash budget are then re-ranked using exact Jaccard similarity on the fingerprint sets.

3.4. HYBRIDSOURCETRACKER (HST)

We conducted additional analysis to assess how combining multiple detection techniques can improve the effectiveness and efficiency of code similarity evaluation, resulting in the development of HST system. The proposed

HST, illustrated in Figure 1, applies two detectors in sequence: (1) SOURCETRACKER powered with QdrantDB to accelerate the first lookup phase, and (2) the Winnowing algorithm.

In the first stage, we use SOURCETRACKER to rapidly select candidates by querying a vector database. This stage utilizes vector search to retrieve the top 100 most similar code snippets based on cosine similarity of embeddings, achieving a logarithmic time complexity with respect to the database size. By limiting subsequent analysis to this reduced candidate set of 100 snippets, we effectively make the second stage operate in constant time, thus maintaining the overall logarithmic complexity. Figure 4 shows how relying on top 100 candidates returned by SOURCETRACKER preserves the effectiveness of the Winnowing algorithm.

In the second stage, we rerank the top 100 list based on the Jaccard similarity of code snippets transformed using the Winnowing algorithm [10]. This approach enables a fine-grained assessment of similarity. By identifying shared fingerprints between the candidate and query snippets, this step enhances the ranking and improves detection efficacy.

This two-tiered approach ensures that coarse similarity detection is performed at high speed, while the final decision leverages a more exact, token-level comparison, yielding a balance between scalability efficiency and retrieval performance.

As a result, users receive a ranked list of “golden code snippets” which constitute the most likely original snippets used by the LLM, possibly subject to adaptation, to produce its output. Generated code can hence be paired with relevant metadata, such as authorship or even licensing information.

4. Results

We present a comprehensive evaluation of HST, which combines syntactically similar code snippet retrieval using SOURCETRACKER with fingerprint-based re-ranking via Winnowing. Our evaluation examines the impact of (1) different dataset sizes (from 1000 to 100 000 samples), (2) varying window lengths (ranging from 7 up to 480 tokens per window, with a visual illustration provided in Figure 3), and (3) different clone types (1 and 2)—and empirically compares the efficiency of the proposed methodologies.

To complement the qualitative illustration in Figure 3, we report in Table 1 a descriptive analysis of the evaluated query windows. From 7 to 480 tokens, mean source coverage increases from 4.99% to 79.63%, the mean




Figure 3: Incremental representation of different window sizes, ranging from 7 to 60 tokens per window; the complete 120-token representation is omitted.

Table 1: Descriptive characteristics of 10 000 unmodified query fragments per nominal token-window size. Continuous measures are reported as mean $\pm$ SD. Full source included is the percentage of snippets fully contained in the window; source coverage is the percentage of each source represented by the window; non-empty lines contain at least one non-whitespace character; and identifier occurrences exclude common keywords.

Window (tokens)	Full source included (%)	Source coverage (%)	Non-empty lines	Identifier occurrences
7	0.29	4.99 ± 10.03	1.48 ± 0.61	2.12 ± 1.20
15	1.40	9.97 ± 16.56	2.24 ± 0.94	4.39 ± 2.05
30	3.84	17.69 ± 23.36	3.70 ± 1.54	8.50 ± 3.56
60	9.03	29.68 ± 30.39	6.45 ± 2.66	16.31 ± 6.53
120	20.69	45.88 ± 35.41	11.38 ± 4.82	30.49 ± 13.07
240	38.72	63.62 ± 35.14	19.41 ± 9.29	53.71 ± 27.02
480	59.04	79.63 ± 28.94	30.56 ± 18.23	86.56 ± 55.24

number of non-empty lines from 1.48 to 30.56, identifier occurrences from 2.12 to 86.56, and the percentage of fully included sources from 0.29% to 59.04%. These results show how larger windows are generally more representative of their source. In practical terms, a 30-token window spans on average about 3.7 non-empty lines and 18% of its source, roughly a function signature with one or two statements, while a 60-token window spans about

6.5 lines and 30% of its source, roughly a small but complete function body. This transition from partial to self-contained code units is a plausible explanation of why both HST and SOURCETRACKER start matching Winnowing at 30 tokens and overtake it from 60 tokens (Section 4.3).

We compare HST against the two baselines described in Section 3: **Winnowing** alone and **SOURCETRACKER**.

4.1. Evaluation metrics

To assess the performance of our provenance retrieval methodologies, we use standard ranked retrieval evaluation measures. These metrics quantify both the system’s ability to return at least one correct match among the top- N candidates (Recall@ N) and the rank of such matches in the result list (MRR).




Figure 4: Comparison of performance between SOURCETRACKER and the Winnowing method, conducted to determine the optimal number of snippets to retrieve in the initial stage of HST, aiming to match Winnowing’s effectiveness. The plot shows retrieval performance at R@100 for SOURCETRACKER and R@1 for Winnowing, averaged across dataset sizes from 1000 to 100 000 samples. Solid lines depict exact match variants, while dashed lines indicate always-replacement variants. Our model (blue) consistently achieves retrieval performance that is equal to or exceeds Winnowing (red). This guarantees that the first-stage retrieval maintains the overall effectiveness of the Winnowing process.

4.1.1. Recall@N

We evaluate retrieval quality using the Recall@N metric, defined in Equation (1) as the fraction of queries for which at least one relevant item is retrieved within the top- N ranked results. Formally, let $Q = \{q_1, q_2, \dots, q_{|Q|}\}$ denote the set of queries. For each query $q_i \in Q$, let G_i be the set of ground-truth relevant items, and let $R_i^{(N)}$ be the set of top- N items retrieved by the system. Then:

$$\text{Recall@N} = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \mathbf{1}(R_i^{(N)} \cap G_i \neq \emptyset), \quad (1)$$

where $\mathbf{1}(\cdot)$ denotes the indicator function, equal to 1 if the argument is true and 0 otherwise. In our context, a higher Recall@N indicates that the system retrieves at least one original snippet within the top- N candidate results.

4.1.2. Mean Reciprocal Rank (MRR)

We also report performance using the Mean Reciprocal Rank (MRR), defined in Equation (2), which evaluates the average ranking position of the first relevant result across all queries. MRR is particularly suitable for provenance retrieval, as it measures how highly the known source is ranked among the retrieved candidates.

Formally, let rank_i denote the position of the first relevant item for query q_i . If no relevant item exists in the retrieved list, we define $\frac{1}{\text{rank}_i} = 0$. The MRR is defined as:

$$\text{MRR} = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \frac{1}{\text{rank}_i}. \quad (2)$$

Using the notation from Equation (1), rank_i can be written as:

$$\text{rank}_i = \min\{k \mid r_{i,k} \in G_i\}, \quad (3)$$

where $r_{i,k}$ is the k -th ranked item retrieved for query q_i . In our domain, a higher MRR indicates that the system tends to rank syntactically similar code snippets higher, thereby reducing manual verification effort.

4.2. Baselines

Figure 5 shows the impact of window size, random replacement, and dataset size on our baselines. While both of our baselines performed poorly




Figure 5: The figures illustrate the Recall@1 ($R@1$) across window sizes ranging from 7 to 480 tokens with three different dimensions of search space—1000, 10 000, and 100 000 samples. We emphasize that both dataset size and window size significantly influence the performance of SOURCETRACKER and the Winnowing algorithm. While SOURCETRACKER performs better with larger windows (up to 240 tokens), the Winnowing algorithm shows greater resilience to variations in both window size and dataset size.

for window sizes below 15 tokens, the Winnowing algorithm appeared more resilient to both a larger search space and a smaller window size, making it more effective. We also emphasize that both methodologies performed on par when comparing Type 1 to Type 2 clones, demonstrating resilience to domain adaptation. SOURCETRACKER performed on par or better than the Winnowing algorithm when dealing with large windows (more than 60 tokens per window), showing how Large Language Models benefit from large contexts.

4.2.1. Time performance analysis

While the MOSS system with the Winnowing algorithm showed superior performance across different datasets and window sizes, using a vectorial DB offers a significant improvement in lookup speed, transitioning from the Winnowing’s linear complexity to logarithmic complexity. Figure 8 displays a clear view of how scaling linearly in time complexity results in a time bottleneck when dealing with larger datasets. Figure 7 demonstrates the practical advantages of our vector database approach. Unlike the Winnowing linear scaling, Qdrant maintains consistent logarithmic performance across dataset sizes. The extrapolated results show that even at 10^{11} samples, query latency remains in the order of milliseconds, making our approach viable for

real-world dataset sizes.

4.3. HST performance

HST, illustrated in Figure 1, was designed to address the challenges of conducting a large-scale retrieval task. The approach aims to ensure efficient retrieval while maintaining the effectiveness of traditional fingerprinting techniques.




Figure 6: MRR (%) comparison across different window sizes between SOURCETRACKER, WinoWing, and HST. Each subplot reports the performance trend as the dataset size increases, while different panels correspond to different window sizes. The proposed HST consistently matches or outperforms the baseline methods for window sizes larger than 30 tokens, while preserving the same logarithmic-time retrieval complexity as SOURCETRACKER, enabled by Qdrant-based indexing.

In Figure 6 and Table 5, we compare the performance of SOURCETRACKER, the WinoWing method [10], and HST, which sequentially applies both techniques. We performed all experiments using the procedure described in Algorithm 1 50% of the time, where frequent words in the fragment were replaced. We used the Mean Reciprocal Rank (MRR) as the metric to report the results, which measures the rank position of the correct item.

The results show that HST achieves performance comparable to WinoWing when the window size is at least 30 tokens (Table 5). Furthermore, for window sizes larger than 60 tokens—which correspond to the values proposed by Ziegler et al. [7] for plagiarism detection—both HST and SOURCE-

Table 2: MRR (%) results using the OLMoTRACE solution are presented for a dataset test size of 100 000, with clone Type 2. While OLMoTRACE demonstrates strong performance with a window size of 7 tokens, it struggles to capture contextual information when larger fragments on adapted code are utilized. Consequently, the model underperforms compared to HST solution (Δ HST).

Window size	MRR	Δ HST
7	26.1	+15.6
15	31.4	-3.6
30	39.2	-29.7
60	44.3	-45.3
120	53.4	-42.0
240	59.1	-39.1
480	67.3	-32.0

TRACKER outperform Winnowing. We also note that HST exhibits more robust performance preservation (Figure 6) for window sizes exceeding 30 tokens, particularly when testing on a larger test set.

Finally, we analyze the computational complexity of our hybrid approach. Theoretically, the first retrieval stage using Qdrant’s HNSW-based vector search should scale logarithmically with respect to dataset size. This is confirmed experimentally in Figure 7. Since we select a constant number of documents (top 100) from the first retrieval stage, the Winnowing algorithm operates on a fixed-size candidate set in constant time, ensuring the overall HST maintains logarithmic complexity with respect to dataset size.

4.4. The OLMoTRACE solution

We tested OLMoTRACE [15] as a baseline for tracing language model outputs back to training data. The system employs a suffix array-based approach, extending the infini-gram framework [22], achieving logarithmic-time search while providing exact-matching guarantees. Its methodology uses maximal span detection for document retrieval and ranking with a BM25 algorithm, making it a scalable and technically efficient solution.

Despite its efficiency advantages, our preliminary experiments (see Table 2) showed that OLMoTRACE was ineffective—when compared with HST—at detecting token windows larger than 15 tokens, showing a gap in terms of MRR (%) starting from -29.7 up to -45.3. The reason lies in its reliance on maximal verbatim spans. Under identifier renaming, every replaced identifier

interrupts the span that would otherwise match the source. Short windows rarely contain a renamed identifier and are matched as a whole, which explains the advantage of OLMOTRACE at 7 tokens; as the window grows, the number of identifier occurrences grows with it (Table 1), the query is fragmented into short exact spans shared by many documents, and the true source no longer ranks first. Fingerprint-based comparison relies on the overlap of fingerprint sets rather than on contiguous matches and is therefore far less sensitive to this fragmentation. This finding is consistent with prior studies highlighting the same limitations [10]. For this reason, we did not include OLMOTRACE in the final set of baselines considered in our evaluation.




Figure 7: Qdrant query latency performance across different collection sizes and top-k values. The plots demonstrate consistent logarithmic scaling behavior with minimal latency increase even at large scales. Extrapolated performance (cross markers) shows the system maintaining sub-second response times even at billions of samples, confirming the theoretical $O(\log N)$ complexity advantage over traditional fingerprinting methods.

4.5. Sensitivity analysis

HST performance is sensitive to a set of hyperparameters that can be modified during deployment, resulting in different trade-offs between effectiveness and efficiency. We therefore analyzed three hyperparameters: (1) the candidate-pool size ($\#Snippets K$), set to 100 in the main experiments, (2) the high-frequency fingerprint threshold, set to 0.5% in the main experiments, and (3) the replacement-probability, initially set to 20%.

Table 3 shows that Stage-1 Recall@ K naturally improves when a larger number of snippets is retrieved. However, the performance of the complete HST pipeline is not particularly sensitive to this hyperparameter, since MRR saturates between 50 and 100 snippets per candidate pool. The efficiency




Figure 8: Single-query latency of the MOSS system (Winnowing) versus dataset size. A fitted power-law model $T(N) = aN^b$ ($a = 1.294 \times 10^{-9}$, $b = 1.0573$) indicates near-linear scaling ($b \approx 1$). Both axes are logarithmic; The fitted model is used to extrapolate latency at $N = 10^9$, 10^{10} , 10^{11} samples (cross markers).

analysis also shows that the absolute difference in re-ranking time is negligible for individual queries, as all the measured values remain below 1 ms. Moreover, the latency of the first stage is essentially unaffected by the number of retrieved candidates in this range (Figure 8 reports nearly identical latency curves for top-k = 1, 10 and 100), so a smaller candidate pool would not translate into a faster pipeline, while a larger one brings no gain in MRR. We therefore keep $K = 100$ as a recall-oriented choice with no measurable cost.

Regarding the high-frequency fingerprint threshold, 0.5% represents an efficiency-oriented operating point. Increasing the threshold from 0.25% to 0.5% produces an improvement of 13.9 MRR percentage points. Higher thresholds further improve retrieval effectiveness, but with progressively smaller MRR gains and higher query costs. Since a higher threshold retains more fingerprints in the index, both effectiveness and query time increase monotonically with it; 0.5% is the efficiency-oriented operating point used in the main experiments, and deployments preferring recall over speed can select a

Table 3: Sensitivity of HST to the candidate-pool size and the high-frequency fingerprint threshold. Results are averaged across all window sizes. Re-ranking time refers only to the Winnowing re-ranking stage and not to the complete end-to-end retrieval latency.

Candidate-pool size			
# Snippets K	Stage-1 Recall@ K (%)	HST MRR (%)	Re-ranking time (ms)
25	74.8	70.8	0.094
50	77.2	71.6	0.164
100	79.4	71.8	0.305
200	81.6	71.7	0.596
High-frequency fingerprint threshold			
Threshold (%)	HST MRR (%)	Recall@100 (%)	Query time (ms)
0.1	32.5	33.6	0.35
0.25	57.9	62.3	1.57
0.5	71.8	80.7	3.11
1.0	73.4	88.1	5.22
2.0	73.8	90.7	7.41

Table 4: Sensitivity to the identifier-replacement probability. All the results are MRR (%) values are averaged across all seven window sizes and five random seeds.

Replacement probability (%)	Winnowing	SourceTracker	HST
10	66.1	64.1	72.6
20	65.1	63.2	71.8
40	63.3	61.2	70.3
60	61.6	59.3	68.6

higher value from Table 3.

Finally, we inspected how the replacement-probability hyperparameter introduced in algorithm 1 impacts the performance of HST, SOURCETRACKER, and the Winnowing algorithm. Table 4 shows that HST is slightly more robust overall than either component alone. When the replacement probability increases from 10% to 60%, average MRR decreases by 4.5 percentage points for Winnowing, 4.8 percentage points for SOURCETRACKER, and 4.0 percentage points for HST.

The results show a monotonic and gradual degradation as the replacement probability increases, rather than a sudden drop at one specific value.

Moreover, HST remains the strongest method under every tested replacement probability.

5. Are errors truly errors?

While we believe that evaluating a retrieval system with classical relevance metrics, such as MRR, is fundamental to understanding the methodology’s true efficacy, our use case requires a parallel, less strict evaluation process. As we mentioned in Section 1, while LLMs might verbatim reproduce part of the training corpus, it is still unclear to what extent [9], and if the generation is attributable to just one source. Our retrieval system is thus designed to retrieve a list of candidate snippets representing plausible provenance matches for the generated code. We decided to evaluate our model not only on ground-truth retrieval but also using an LLM as a judge (Qwen3 Coder [23]), and two LLMs (Codestral 25.08 [24] and deepseek-v4-pro [25]) to measure the inter-judge agreement (see Section 5.1). This judge assesses the similarity between the user’s query and the retrieved code snippets. These snippets are not part of the ground truth, but will be presented to the end user as potential sources of inspiration for the input snippet. This assessment uses a scale from 1 to 5, where 1 indicates that the fragment is completely independent of the snippet, and 5 signifies that the fragment is clearly derived from the snippet. Figure 9 illustrates that even when the model fails to retrieve the “golden snippet”, there is still a noticeable level of similarity when examining large fragments of code—suggesting that the snippet may have influenced the generation process. We highlight two main factors that cause the model to retrieve erroneous snippets and fallacious suggestions: 1) small fragment windows make the retrieval process much more difficult for HST, and 2) the adaptation of the code snippets—in our case, the frequent word replacement (Clones Type 2)—deteriorates the retrieval process.

By combining the results from clone Types 1 and 2 (see the right image of Figure 9), we experienced ascending average ratings from 1.3 to 2.6 as the context window size gets larger (see Table 7). We highlight how, with a minimum of 30 tokens per window, the model starts retrieving relevant results more effectively, even when they are not part of the ground truth.

We conducted a permutation ANOVA (10 000 permutations) to assess whether different window sizes yielded distinct results. For Clones Type 1, the analysis produced an F-statistic of 26.6 with a p-value of < 0.001 , and

Table 5: Comparison of MRR (%) outcomes using 100 000 fragments across a range of window sizes. Starting from a window size of 30 tokens, HST consistently achieves higher MRR than both the Winnowing baseline and SOURCETRACKER across all evaluated configurations.

Methodology	7	15	30	60	120	240	480
SOURCETRACKER	5.5	21.9	53.8	81.2	90.5	96.2	98.8
Winnowing	20.4	46.0	68.3	84.2	91.3	95.4	97.8
HST	10.5	35.0	68.9	89.6	95.4	98.2	99.3

for Clones Type 2, an F-statistic of 14.2 with a p-value also < 0.001 . These findings indicate that window size has a significant effect on the outcomes.

We evaluated the ranking mechanism of the LLM judge by having it assess the similarities between code snippets from the ground truth (true positive pairs), where ideally every couple of fragment-source should get five points out of five. Our results (see Table 7) indicate that an LLM can identify whether a code snippet comes from its training set; however, when the task is framed with a context as short as 7 tokens per window, it becomes challenging to clearly determine whether the snippet originates from a particular source. We applied a permutation ANOVA to assess whether varying the window size produced different outcomes. For Clones Type 1, this yielded an F-statistic of 64.0 with a p-value < 0.001 , and for Clones Type 2, an F-statistic of 51.4, with a p-value also < 0.001 . We observed that the length of the context window has a small but significant effect. Specifically, when the true context spanned 7 to 30 tokens per window, it was more often assigned a score of 4 out of 5, resulting in a less reliable assessment.

A lower degree of similarity between the erroneously coupled fragment-source can therefore result not only from a retrieval error, but may also be attributed to the LLM acting as evaluator.

5.1. Inter-judge agreement

To verify that the results coming from `qwen3-coder` were not dependent on a single LLM judge, we evaluated the agreement using two code-oriented models: `codestral-2508` and `deepseek-v4-pro`. The three models evaluated the same query-retrieved snippet pairs using the same five-point similarity scale.

We measured their agreement using Krippendorff’s ordinal alpha (where 1 indicates perfect agreement, 0 indicates agreement no better than chance,




Figure 9: Results of the LLM error similarity evaluation: The left figure illustrates how different clone type errors (verbatim clone Type 1 and frequent word replacement clone Type 2) result in distinct suggestions, leading to worse recommendations from our system when fragments are adapted. Spearman rank correlations revealed significant monotonic relationships between context window size and similarity scores (Clone Type 1: $r_s = 0.446$, $p < 0.001$; Clone Type 2: $r_s = 0.360$, $p < 0.001$). The second figure shows the similarity error results aggregated by clone type. Both figures indicate that larger context windows in generation lead to more inherent snippets retrieved.

and negative values indicate disagreement beyond chance), exact agreement (the percentage of examples for which all judges assigned the same score), agreement within one point (the percentage of examples for which all three judges’ scores differed by no more than one point), and pairwise Spearman correlations (ranging from -1 to 1 , where 1 indicates the same ranking, 0 no monotonic association, and -1 the opposite ranking).

As shown in Table 6, the three judges obtained an ordinal alpha of 0.651 across $1,262$ jointly evaluated pairs. Exact agreement was 43.98% , while in 82.57% of the cases the three scores differed by at most one point. Pairwise Spearman correlations ranged from 0.670 to 0.765 , showing that the models were generally consistent in ranking the similarity of the retrieved snippets.

The models showed some differences in their use of the scale. Qwen3 Coder assigned the highest average score (2.005), followed by DeepSeek (1.748) and Codestral (1.620). However, all three judges maintained the same overall trend, with similarity scores increasing for longer context windows. Therefore, the main conclusion of the LLM-based evaluation is not dependent on one specific LLM judge.

Table 6: Agreement among the three LLM-based similarity judges aggregated across all window sizes and clone types.

Agreement measure	Result
Krippendorff’s ordinal α	0.651
Exact agreement	43.98%
Agreement within one point	82.57%
Pairwise Spearman ρ	0.670–0.765
Jointly evaluated pairs	1,262

6. Discussion

We designed HST to address the scalability challenges of provenance retrieval over large-scale training datasets for generative LLMs, thereby supporting downstream authorship attribution and plagiarism assessment. Our results highlight three main insights: (1) the classic Winnowing algorithm is still one of the most effective methodologies, but it is suboptimal in terms of time complexity when applied to large-scale datasets. (2) Fine-tuned encoders improve retrieval effectiveness with larger fragments and enable the usage of HNSW [18], leading to queries with logarithmic time complexities w.r.t. dataset size. (3) The sequential integration of both approaches obtains the best results in terms of both efficacy and efficiency when dealing with ≥ 30 token window sizes.

When we compare SOURCETRACKER with Winnowing [10] in isolation, Winnowing consistently maintains higher robustness across small window sizes and large datasets. This outcome aligns with its well-established reliability as a fingerprinting technique that does not rely on semantic embeddings. However, we observed that SOURCETRACKER performs particularly well when sequences exceed 120 tokens, confirming that encoder-based methods benefit from longer contextual windows that capture richer semantics.

We addressed scalability requirements by shifting the methodology. Even though Winnowing performs strongly in terms of detection, its linear complexity becomes a bottleneck when applied at the scale of billions of training examples. By integrating a vector database in the first stage, we reduced the search complexity to logarithmic, which substantially accelerates the retrieval process.

For window sizes larger than 60 tokens—the interest threshold suggested in prior work for reliable plagiarism detection [7]—we achieved superior per-

Table 7: Comparison of LLM judge similarity score statistics for different window sizes between Ground Truth and HST: mean, standard deviation (σ), and median, with clone Types 1 and 2 results averaged.

Window size	Ground Truth		HST	
	Mean $\pm \sigma$	Median	Mean $\pm \sigma$	Median
7	3.47 ± 1.15	4	1.36 ± 0.66	1
15	4.03 ± 0.76	4	1.56 ± 0.85	1
30	4.36 ± 0.56	4	2.01 ± 1.05	2
60	4.47 ± 0.60	5	2.20 ± 1.10	2
120	4.69 ± 0.49	5	2.30 ± 1.07	2
240	4.79 ± 0.47	5	2.51 ± 1.07	3
480	4.91 ± 0.28	5	2.63 ± 1.15	3

formance compared to Winnowing alone, showing how our two-stage process is capable of maintaining the representation capabilities of both Winnowing and SOURCETRACKER.

Additionally, we observed that simple obfuscation strategies (Clone Type 2 fragments), such as identifier replacement, caused noticeable performance degradation primarily in code snippets outside the ground truth, thereby impairing provenance retrieval and, consequently, limiting downstream authorship attribution. This finding highlights the importance of our double evaluation approach—such an effect would have remained unnoticed under a conventional ground truth-only evaluation. As it seems to be obvious, the more the LLM generates “creative” code, the farther it is from the original snippets of code, the more difficult it is to identify its provenance.

We thus want to emphasize the need for broader benchmarking evaluations of these systems. Both old-fashioned ground truth evaluation and the LLMs judge paradigm, when used alone, cannot represent the efficacy of the methodology. Obtaining the ground truth result is as essential as having a herd of good related suggestions, as LLMs’ generation process remains a black box to this day.

Finally, we underscore the main limitation of this approach—its applicability is contingent on the availability of a suitable training set—and frame this as part of a broader call for transparency. While generative methods that explicitly encode authorship information may offer a practical path forward, relying exclusively on the internal memorization tendencies of LLMs is fundamentally precarious, as the degree and the underlying mechanism of such

memorization are still poorly understood [5]. We therefore urge greater transparency in the deployment of Large Language Models, making provenance tracking—as an enabler for authorship attribution—a central requirement and treating the release of training datasets and their underlying sources as a *de facto* requirement.

7. Related work

We focus on provenance tracking in the context of LLM-generated code, as multiple studies have shown that large language models can memorize and reproduce training data [5, 9, 7]. Such behavior, commonly referred to as “recitation” [7], may result in verbatim or adapted reuse of training examples, raising concerns related to plagiarism and copyright infringement [1, 2, 3]. Identifying the provenance of generated code can therefore provide useful evidence for downstream tasks such as authorship attribution, plagiarism assessment, and license-compliance analysis.

In this section, we review prior work on memorization and plagiarism detection for natural language (Section 7.1), code plagiarism and clone detection (Section 7.2), and related approaches that provide methodological foundations for scalable provenance retrieval. Finally, we discuss the legal and ethical implications of these problems (Section 7.4).

7.1. Memorization and Plagiarism Detection in Natural Language

Lee et al. [9] demonstrated that LLMs plagiarize in three distinct ways: verbatim copying, paraphrasing, and by repeating ideas present in the training data. They also demonstrated how larger models tend to memorize more of the training dataset. Their work on the PAN Dataset achieved detection accuracies of 0.92 for non-plagiarism cases, 1.0 for verbatim plagiarism, 0.88 for paraphrasing, and 0.62 for idea-based plagiarism. They employed a multi-step classification approach using bag-of-words search for candidate retrieval followed by text decomposition and similarity computation.

Carlini et al. [5] further analyzed how LLMs memorize training data, providing empirical evidence that memorization increases with (1) model capacity, (2) data duplication frequency, and (3) context length—a phenomenon they termed “discoverability”, where some memorization only becomes apparent under the condition of giving the LLM an input that reveals the memorized text.

Lee et al. [26] also introduced LLMs as judges capable of binary classifying whether or not a snippet of code is plagiarized from another. They developed the PlagBench corpus, consisting of 46.5K text pairs covering three types of LLM-aware plagiarism (verbatim, paraphrase, and summary cases). The dataset was generated using GPT-3.5 Turbo and GPT-4 Turbo. Their results demonstrated that modern LLMs, such as Llama3 and GPT-4, can outperform existing plagiarism checkers trained explicitly for the task, achieving binary classification accuracies above 80% through few-shot chain-of-thought prompting—these and similar results constitute our main motivation for also experimenting with an LLM-based judge.

Panaitecsu-Liess et al. [27] demonstrated that models with low perplexity scores for specific text segments are more likely to have memorized and replicated that content. They extended the memorization definition from Carlini et al. [5] by using perplexity as a feature for classification, following the intuition that if the perplexity is too low—and thus the model produces output tokens with high confidence—there is a high chance that the sentence has been memorized.

Our study differs from earlier work in both its scope and application domain. Rather than auditing whether an LLM has memorized particular training examples or directly classifying generated code as plagiarized, we focus on scalable provenance retrieval. Given a code fragment and a known training corpus, our objective is to retrieve the candidate source snippets most closely associated with that fragment. The resulting provenance information can subsequently support authorship attribution, when author metadata is available, as well as downstream plagiarism and license-compliance assessments.

Since modern training corpora contain billions of code snippets, scalability constitutes a central requirement of this retrieval task. To address this challenge, HST combines HNSW-based vector similarity search [18] with fingerprint-based re-ranking, enabling efficient retrieval over very large code collections.

7.2. Code plagiarism detection and clone detection

In the realm of code plagiarism detection, fingerprinting techniques have proven particularly effective. The Winnowing algorithm [10], which forms the foundation of the MOSS system [11], has been broadly adopted by academic institutions and employed in investigations of intellectual-property infringement. Winnowing operates by selecting representative hashes—termed

fingerprints—that remain robust to minor syntactic or formatting modifications.

Henderson et al. [11] demonstrated using the winnowing-based MOSS system that code generation models can generate function implementations that substantially overlap with reference implementations from training data. These algorithms typically operate in linear time relative to the size of the inspected dataset. However, LLMs trained on extensive collections of code snippets have broadened the search space considerably, making large-scale provenance retrieval increasingly challenging. Recent pre-training datasets far exceed the size of billions of samples [12, 14, 13], making algorithms with linear time complexities inadequate for modern-scale applications.

Recent work by Bifulco et al. [28] investigated the quality of provenance links provided by modern code generation systems like Gemini and Bing Copilot. Their analysis, using code clone detection tools (including Code2vec [29]) calibrated with a 20-token minimum threshold to detect Type 1 and Type 2 clones, revealed that both models provide a mix of relevant and non-relevant links, with Gemini providing up to 48% unrelated links. They conjectured that these models rely on search engines for similarities, but found the overlap between search engine results and provided links to be low. Importantly, their findings indicated that clone detection and textual similarity tools produced results comparable to manual analysis, validating their utility for assessing code similarity.

Liu et al. [15] recently addressed the task with a focus on efficiency. Using the infinity-gram system [22], their method achieved logarithmic query time complexity by relying on verbatim matches. Their approach is notable for being the first to address the challenge with large corpus; in their production system, they can trace an LLM output against 4.6 trillion training tokens, making it the first to analyze such a data scale.

We rely on both vectorial and fingerprinting techniques for our hybrid system, aiming to achieve the efficiency demonstrated by Liu et al. [15] with the HNSW [18] algorithm, while overcoming problems related to the sensitivity of verbatim lookups to noise [10]. In our experimentation with large generated outputs, we combine vectorial lookups with the Winnowing technique to refine our research. This way, we achieved logarithmic time complexity for search, as previously done by Liu et al. [15], while retaining the effectiveness of traditional fingerprinting techniques [10].

7.3. Semantic code understanding and representation learning

Ahmad and Zhang [30] introduced a Python-oriented API framework that uses static analysis to extract repository-level information, including function definitions, dependencies, metadata, cross-file relationships, and semantic context. Although designed to improve LLM interactions rather than provenance tracking, it demonstrates how project-specific information can reveal code relationships beyond lexical similarity. In complementary work, Ahmad and Zhang [31] proposed a dual-encoder architecture that separately processes source-code tokens and AST nodes and combines their representations through cross-attention. While developed for code optimization, this architecture illustrates how lexical and structural information can be integrated to improve semantic code understanding.

These works complement HST but address different objectives. The framework in [30] enriches the repository context available to an LLM, while [31] improves the representation of individual code fragments. In contrast, HST focuses on scalable provenance retrieval: SOURCETRACKER generates dense representations for approximate nearest-neighbor search, followed by Winnowing-based re-ranking. Nevertheless, repository-level context and joint token-AST representations provide promising directions for improving SOURCETRACKER’s ability to retrieve more extensively modified Type 2 clones.

7.4. Legal and ethical considerations

Henderson et al. [11] examined the legal landscape surrounding foundation models and copyright. They note that, since most online content is protected by copyright at creation, using such material for training could constitute infringement. The fair use doctrine permits the public to use copyrighted material for certain purposes, especially when the resulting work is transformative. As Lemley and Casey [32] argued, training a machine learning model on copyrighted data is likely considered fair use when the final model does not directly generate content. However, when training and deploying foundation models for generative use cases, the analysis becomes more complex, as these models can generate content similar to copyrighted material, potentially affecting the original data creators. This issue arises from the LLMs’ ability to memorize copyrighted material in their weights [33] implicitly.

Several tests exist to examine nonliteral infringement, such as the

Abstraction-Filtration-Comparison test and the Structure, Sequence, and Organization (SSO) test. As noted by Henderson et al. [11], in software

copyright cases, proving nonliteral infringement is challenging, largely because courts recognize that nonliteral elements, such as algorithms, are not within the scope of protection that copyright law provides. This complexity is further amplified by LLMs’ ability to generate code that may substantially overlap with reference implementations.

7.5. From provenance to attribution in practice

Two practical requirements follow from our design: access to the training corpus, and author and license metadata for every retrieved snippet. For openly released training sets, both are already met by existing infrastructure. THESTACKV2 [12] is derived from the Software Heritage archive [34], and each file carries the Software Heritage persistent identifiers (SWHIDs) of its content, directory, revision and snapshot, together with the repository name, branch, commit dates, and the SPDX licenses detected by ScanCode. Since SWHIDs are persistent, content-addressed identifiers that resolve to the full version history of the archive, each golden snippet returned by HST can be mapped to its origin repository, its commits and hence its authors, with no additional data collection. The CodeCommons initiative [35], hosted by Software Heritage [34], builds on the same infrastructure precisely to attach provenance, licensing and attribution information to code used for training language models. Operationalising HST for such a model therefore amounts to a one-time indexing of its training set (encoding and HNSW index construction) and, at query time, returning SWHIDs alongside the retrieved snippets. OLMOTRACE [15] follows the same deployment pattern for models trained on openly released data, which suggests that the approach is viable in production settings.

For models whose training data is not disclosed, HST cannot be applied, and the same holds for any post-hoc attribution method. This is why we regard the disclosure of training data, or at least of a queryable index of it, as a prerequisite for responsible deployment rather than as an optional courtesy. Regulation is moving in this direction: under Article 53 of the EU AI Act [36], providers of general-purpose AI models must make publicly available a sufficiently detailed summary of the content used for training, following the template adopted by the European Commission in July 2025 [37]. A summary is not a searchable corpus, but it establishes the principle that training data provenance is a matter of public interest, and tools such as HST show what becomes possible once that principle is taken one step further.

8. Conclusion

Providing proper attribution across LLM-based code generation remains a significant challenge. Our work addresses this issue by retrieving the code snippets most likely to have influenced the model’s output. Specifically, we design HST, a hybrid system that, after the generation phase, performs a similarity search over the training corpus to identify and present to the end user the most similar code snippets, along with associated metadata such as authorship information. This approach directly addresses the challenges posed by LLMs: we combined the efficiency of vector lookups to manage the vast scale of training data with the effectiveness of fingerprint techniques. Our fine-tuned model, SOURCETRACKER, is available along with the training code and evaluation experiments (see Section 8.1.4). We propose a methodology for fairly attributing source code to its authors. However, we must acknowledge a significant limitation: the training dataset needs to be accessible. We conclude by calling for transparency, emphasizing that the training data should be openly available.

8.1. Future work

While our results demonstrate the effectiveness of HST on a 10M-snippet dataset, with evaluations on search spaces up to 100k documents, several research directions remain open.

8.1.1. Beyond Type 1 and Type 2 clones

Our evaluation focuses on Type 1 (verbatim) and Type 2 (identifier-renaming) clones, motivated by recitation and domain-adaptation phenomena in LLM-generated code. Extending both training and evaluation to harder categories, such as Type 3 near-miss clones (with statement insertions, deletions, or reordering) and, where feasible, Type 4 semantic clones, would clarify how far a provenance signal survives under increasingly substantial transformations beyond our current frequent-word replacement scheme.

8.1.2. Adaptive window sizing and context-aware detection

Developing adaptive window-sizing strategies that adjust fragment length based on code characteristics (e.g., function complexity, syntactic density, or language-specific idioms) could improve robustness across diverse coding styles. An additional avenue is to replace fixed-token windows with semantically meaningful units, such as function bodies, classes, or basic blocks,

potentially by integrating abstract syntax tree (AST) analysis or other structural program representations into the fragmentation process. Dual-encoder architectures that process token sequences and AST nodes separately and fuse them through cross-attention [38] are a natural candidate for making SOURCETRACKER more robust to the structural edits that characterise Type 3 clones.

8.1.3. Improved handling of semantic similarity

The LLM-based judge evaluation indicates that many retrieved snippets, although not ground-truth matches, still exhibit non-trivial similarity to query fragments. This motivates retrieval objectives beyond exact ground-truth recall, including ranking by graded similarity or semantic relatedness and learning scoring functions that better align with human judgments, for example through human-in-the-loop or preference-learning approaches. Furthermore, fine-tuning large language models specifically for code-similarity assessment, rather than relying on generic prompting, could yield more reliable semantic comparators to complement HST. Frameworks that extract semantic information from code repositories through static analysis, such as API usage and project-specific structure [30], could likewise provide retrieval signals beyond lexical matching, which would be needed to approach Type 4 clones.

8.1.4. Standardization and community benchmarks

To facilitate reproducible research and enable fair comparison across methods, standardized benchmarks for code provenance tracking are needed. Such benchmarks should include multi-language datasets, controlled levels of code modification (Type 1 through Type 4 clones), and realistic LLM-generated samples, together with agreed-upon protocols for evaluation at scale. Establishing community-driven shared tasks around these benchmarks would likely accelerate methodological progress and clarify the trade-offs between efficiency, robustness, and semantic fidelity.

Data availability

The trained SOURCETRACKER weights, along with the full replication package for all experiments, are publicly available on Zenodo following this link: <https://doi.org/10.5281/zenodo.18375484>. However, note that obtaining THESTACKV2 requires agreeing to the terms set by Software Heritage

and INRIA, as outlined in the dataset license. For a comprehensive overview of the replication package and its components, refer to the README.md file included in the Zenodo repository.

Acknowledgments

This work was made possible by Software Heritage [34], the universal source code archive <https://www.softwareheritage.org>. In particular, we thank the Code Commons project <https://codecommons.org>, for setting up the broader challenge of providing provenance tracking, authorship attribution, and open source licensing information across LLM code generation. Davide D’Ascenzo was financially supported by the Italian National PhD Program in Artificial Intelligence (DM 351 intervento M4C1 - Inv. 4.1 - Ricerca PNRR), funded by NextGenerationEU (EU-NGEU). This work has been partially supported by project SERICS (PE00000014) under the NRRP MUR program funded by the EU-NGEU.

References

- [1] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. de Oliveira Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, A. Ray, R. Puri, G. Krueger, M. Petrov, H. Khlaaf, G. Sastry, P. Mishkin, B. Chan, S. Gray, N. Ryder, M. Pavlov, A. Power, L. Kaiser, M. Bavarian, C. Winter, P. Tillet, F. P. Such, D. Cummings, M. Plappert, F. Chantzis, E. Barnes, A. Herbert-Voss, W. H. Guss, A. Nichol, A. Paino, N. Tezak, J. Tang, I. Babuschkin, S. Balaji, S. Jain, W. Saunders, C. Hesse, A. N. Carr, J. Leike, J. Achiam, V. Misra, E. Morikawa, A. Radford, M. Knight, M. Brundage, M. Murati, K. Mayer, P. Welinder, B. McGrew, D. Amodei, S. McCandlish, I. Sutskever, W. Zaremba, [Evaluating large language models trained on code](#), CoRR abs/2107.03374 (2021). [arXiv:2107.03374](#).
URL <https://arxiv.org/abs/2107.03374>
- [2] B. Rozière, J. Gehring, F. Gloeckle, S. Sootla, I. Gat, X. E. Tan, Y. Adi, J. Liu, T. Remez, J. Rapin, A. Kozhevnikov, I. Evtimov, J. Bitton, M. Bhatt, C. Canton-Ferrer, A. Grattafiori, W. Xiong, A. Défossez, J. Copet, F. Azhar, H. Touvron, L. Martin, N. Usunier, T. Scialom, G. Synnaeve, [Code llama: Open foundation models for code](#), CoRR

abs/2308.12950 (2023). [arXiv:2308.12950](https://arxiv.org/abs/2308.12950), [doi:10.48550/ARXIV.2308.12950](https://doi.org/10.48550/ARXIV.2308.12950).

URL <https://doi.org/10.48550/arXiv.2308.12950>

- [3] R. Li, L. B. Allal, Y. Zi, N. Muennighoff, D. Kocetkov, C. Mou, M. Marone, C. Akiki, J. Li, J. Chim, Q. Liu, E. Zheltonozhskii, T. Y. Zhuo, T. Wang, O. Dehaene, M. Davaadorj, J. Lamy-Poirier, J. Monteiro, O. Shliazhko, N. Gontier, N. Meade, A. Zebaze, M. Yee, L. K. Umapathi, J. Zhu, B. Lipkin, M. Oblokulov, Z. Wang, R. M. V, J. T. Stillerman, S. S. Patel, D. Abulkhanov, M. Zocca, M. Dey, Z. Zhang, N. Fahmy, U. Bhattacharyya, W. Yu, S. Singh, S. Luccioni, P. Villegas, M. Kunakov, F. Zhdanov, M. Romero, T. Lee, N. Timor, J. Ding, C. Schlesinger, H. Schoelkopf, J. Ebert, T. Dao, M. Mishra, A. Gu, J. Robinson, C. J. Anderson, B. Dolan-Gavitt, D. Contractor, S. Reddy, D. Fried, D. Bahdanau, Y. Jernite, C. M. Ferrandis, S. Hughes, T. Wolf, A. Guha, L. von Werra, H. de Vries, [Starcoder: may the source be with you!](#), Trans. Mach. Learn. Res. 2023 (2023).

URL <https://openreview.net/forum?id=KoF0g41haE>

- [4] X. Hou, Y. Zhao, Y. Liu, Z. Yang, K. Wang, L. Li, X. Luo, D. Lo, J. Grundy, H. Wang, [Large language models for software engineering: A systematic literature review](#), ACM Trans. Softw. Eng. Methodol. 33 (8) (2024) 220:1–220:79. [doi:10.1145/3695988](https://doi.org/10.1145/3695988).

URL <https://doi.org/10.1145/3695988>

- [5] N. Carlini, D. Ippolito, M. Jagielski, K. Lee, F. Tramèr, C. Zhang, [Quantifying memorization across neural language models](#), in: The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023, OpenReview.net, 2023.

URL https://openreview.net/forum?id=TatRHT_1cK

- [6] A. M. E. T. Ali, H. M. D. Abdulla, V. Snásel, [Overview and comparison of plagiarism detection tools](#), in: V. Snásel, J. Pokorný, K. Richta (Eds.), Proceedings of the DATESO 2011: Annual International Workshop on DATABASES, TEXTS, SPECIFICATIONS and OBJECTS, Pisek, Czech Republic, April 20, 2011, Vol. 706 of CEUR Workshop Proceedings, CEUR-WS.org, 2011, pp. 161–172.

URL <https://ceur-ws.org/Vol-706/poster22.pdf>

- [7] A. Ziegler, GitHub Copilot research recitation — parrot or crow? a first look at rote learning in GitHub Copilot suggestions, <https://github.blog/ai-and-ml/github-copilot/github-copilot-research-recitation/>, accessed: 2024-09-25 (2021).
- [8] A. Almufarreh, A. Ahmad, M. Arshad, C. W. Onn, R. Elechi, Ethical implications of chatgpt and other large language models in academia, *Frontiers in Artificial Intelligence* 8 (2025) 1615761.
- [9] J. Lee, T. Le, J. Chen, D. Lee, [Do language models plagiarize?](#), in: Y. Ding, J. Tang, J. F. Sequeda, L. Aroyo, C. Castillo, G. Houben (Eds.), *Proceedings of the ACM Web Conference 2023, WWW 2023, Austin, TX, USA, 30 April 2023 - 4 May 2023*, ACM, 2023, pp. 3637–3647. doi:10.1145/3543507.3583199. URL <https://doi.org/10.1145/3543507.3583199>
- [10] S. Schleimer, D. S. Wilkerson, A. Aiken, [Winnowing: Local algorithms for document fingerprinting](#), in: A. Y. Halevy, Z. G. Ives, A. Doan (Eds.), *Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data, San Diego, California, USA, June 9-12, 2003*, ACM, 2003, pp. 76–85. doi:10.1145/872757.872770. URL <https://doi.org/10.1145/872757.872770>
- [11] P. Henderson, X. Li, D. Jurafsky, T. Hashimoto, M. A. Lemley, P. Liang, [Foundation models and fair use](#), *J. Mach. Learn. Res.* 24 (2023) 400:1–400:79. URL <http://jmlr.org/papers/v24/23-0569.html>
- [12] A. Lozhkov, R. Li, L. B. Allal, F. Cassano, J. Lamy-Poirier, N. Tazi, A. Tang, D. Pykhtar, J. Liu, Y. Wei, T. Liu, M. Tian, D. Kocetkov, A. Zucker, Y. Belkada, Z. Wang, Q. Liu, D. Abulkhanov, I. Paul, Z. Li, W. Li, M. Risdal, J. Li, J. Zhu, T. Y. Zhuo, E. Zheltonozhskii, N. O. O. Dade, W. Yu, L. Krauß, N. Jain, Y. Su, X. He, M. Dey, E. Abati, Y. Chai, N. Muennighoff, X. Tang, M. Oblokulov, C. Akiki, M. Marone, C. Mou, M. Mishra, A. Gu, B. Hui, T. Dao, A. Zebaze, O. Dehaene, N. Patry, C. Xu, J. J. McAuley, H. Hu, T. Scholak, S. Paquet, J. Robinson, C. J. Anderson, N. Chapados, et al., [Starcoder 2 and the stack v2: The next generation](#), *CoRR* abs/2402.19173 (2024). arXiv:2402.19173, doi:10.48550/ARXIV.2402.19173. URL <https://doi.org/10.48550/arXiv.2402.19173>

- [13] G. Penedo, H. Kydlíček, L. B. Allal, A. Lozhkov, M. Mitchell, C. A. Raffel, L. von Werra, T. Wolf, [The fineweb datasets: Decanting the web for the finest text data at scale](#), in: A. Globersons, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. M. Tomczak, C. Zhang (Eds.), Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024, 2024.
URL http://papers.nips.cc/paper_files/paper/2024/hash/370df50ccfd8bde18f8f9c2d9151bda-Abstract-Datasets_and_Benchmarks_Track.html
- [14] T. Nguyen, C. V. Nguyen, V. D. Lai, H. Man, N. T. Ngo, F. Deroncourt, R. A. Rossi, T. H. Nguyen, [CulturaX: A cleaned, enormous, and multilingual dataset for large language models in 167 languages](#), in: N. Calzolari, M.-Y. Kan, V. Hoste, A. Lenci, S. Sakti, N. Xue (Eds.), Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024), ELRA and ICCL, Torino, Italia, 2024, pp. 4226–4237.
URL <https://aclanthology.org/2024.lrec-main.377/>
- [15] J. Liu, T. Blanton, Y. Elazar, S. Min, Y. Chen, A. Chheda-Kothary, H. Tran, B. Bischoff, E. Marsh, M. Schmitz, C. Trier, A. Sarnat, J. James, J. Borchardt, B. Kuehl, E. Y. Cheng, K. Farley, S. Sreeram, T. Anderson, D. Albright, C. Schoenick, L. Soldaini, D. Groeneveld, R. Y. Pang, P. W. Koh, N. A. Smith, S. Lebrecht, Y. Choi, H. Hajishirzi, A. Farhadi, J. Dodge, [Olmotrace: Tracing language model outputs back to trillions of training tokens](#), CoRR abs/2504.07096 (2025).
[arXiv:2504.07096](#), [doi:10.48550/ARXIV.2504.07096](#).
URL <https://doi.org/10.48550/arXiv.2504.07096>
- [16] C. K. Roy, J. R. Cordy, A survey on software clone detection research, Queen’s School of computing TR 541 (115) (2007) 64–68.
- [17] J. Svajlenko, C. K. Roy, [Evaluating clone detection tools with big-clonebench](#), in: R. Koschke, J. Krinke, M. P. Robillard (Eds.), 2015 IEEE International Conference on Software Maintenance and Evolution, ICSME 2015, Bremen, Germany, September 29 - October 1, 2015, IEEE Computer Society, 2015, pp. 131–140. [doi:10.1109/ICSM.2015.](#)

7332459.

URL <https://doi.org/10.1109/ICSM.2015.7332459>

- [18] Y. A. Malkov, D. A. Yashunin, [Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs](#), IEEE Trans. Pattern Anal. Mach. Intell. 42 (4) (2020) 824–836. [doi:10.1109/TPAMI.2018.2889473](#).
URL <https://doi.org/10.1109/TPAMI.2018.2889473>
- [19] A. Gurioli, F. Pennino, J. Monteiro, M. Gabbrielli, [Mose: Hierarchical self-distillation enhances early layer embeddings](#), in: S. Koenig, C. Jenkins, M. E. Taylor (Eds.), Fortieth AAAI Conference on Artificial Intelligence, Thirty-Eighth Conference on Innovative Applications of Artificial Intelligence, Sixteenth Symposium on Educational Advances in Artificial Intelligence, AAAI 2026, Singapore, January 20-27, 2026, AAAI Press, 2026, pp. 30897–30906. [doi:10.1609/AAAI.V40I37.40348](#).
URL <https://doi.org/10.1609/aaai.v40i37.40348>
- [20] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, I. Sutskever, [Learning transferable visual models from natural language supervision](#), in: M. Meila, T. Zhang (Eds.), Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event, Vol. 139 of Proceedings of Machine Learning Research, PMLR, 2021, pp. 8748–8763.
URL <http://proceedings.mlr.press/v139/radford21a.html>
- [21] Qdrant Team, [Qdrant: Vector search engine](#) (2025).
URL <https://qdrant.tech>
- [22] J. Liu, S. Min, L. Zettlemoyer, Y. Choi, H. Hajishirzi, [Infini-gram: Scaling unbounded n-gram language models to a trillion tokens](#) (2025). [arXiv:2401.17377](#).
URL <https://arxiv.org/abs/2401.17377>
- [23] A. Yang, A. Li, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Gao, C. Huang, C. Lv, C. Zheng, D. Liu, F. Zhou, F. Huang, F. Hu, H. Ge, H. Wei, H. Lin, J. Tang, J. Yang, J. Tu, J. Zhang, J. Yang, J. Yang, J. Zhou, J. Lin, K. Dang, K. Bao, K. Yang, L. Yu, L. Deng, M. Li, M. Xue, M. Li, P. Zhang, P. Wang, Q. Zhu, R. Men, R. Gao,

- S. Liu, S. Luo, T. Li, T. Tang, W. Yin, X. Ren, X. Wang, X. Zhang, X. Ren, Y. Fan, Y. Su, Y. Zhang, Y. Zhang, Y. Wan, Y. Liu, Z. Wang, Z. Cui, Z. Zhang, Z. Zhou, Z. Qiu, [Qwen3 technical report](#), CoRR abs/2505.09388 (2025). [arXiv:2505.09388](#), [doi:10.48550/ARXIV.2505.09388](#).
URL <https://doi.org/10.48550/arXiv.2505.09388>
- [24] Mistral AI, [Codestral v25.08](#), Mistral AI model card, model identifier: codestral-2508 (Jul. 2025).
URL <https://docs.mistral.ai/models/model-cards/codestral-25-08>
- [25] DeepSeek-AI, [Deepseek-v4: Towards highly efficient million-token context intelligence](#), CoRR abs/2606.19348 (2026). [arXiv:2606.19348](#), [doi:10.48550/ARXIV.2606.19348](#).
URL <https://doi.org/10.48550/arXiv.2606.19348>
- [26] J. Lee, T. Agrawal, A. Uchendu, T. Le, J. Chen, D. Lee, [Plagbench: Exploring the duality of large language models in plagiarism generation and detection](#), CoRR abs/2406.16288 (2024). [arXiv:2406.16288](#), [doi:10.48550/ARXIV.2406.16288](#).
URL <https://doi.org/10.48550/arXiv.2406.16288>
- [27] M. Panaitescu-Liess, Z. Che, B. An, Y. Xu, P. Pathmanathan, S. Chakraborty, S. Zhu, T. Goldstein, F. Huang, [Can watermarking large language models prevent copyrighted text generation and hide training data?](#), CoRR abs/2407.17417 (2024). [arXiv:2407.17417](#), [doi:10.48550/ARXIV.2407.17417](#).
URL <https://doi.org/10.48550/arXiv.2407.17417>
- [28] D. Bifulco, P. Cassieri, G. Scanniello, M. D. Penta, F. Zampetti, [Do llms provide links to code similar to what they generate? A study with gemini and bing copilot](#), in: 22nd IEEE/ACM International Conference on Mining Software Repositories, MSR@ICSE 2025, Ottawa, ON, Canada, April 28-29, 2025, IEEE, 2025, pp. 223–235. [doi:10.1109/MSR66628.2025.00042](#).
URL <https://doi.org/10.1109/MSR66628.2025.00042>
- [29] U. Alon, M. Zilberstein, O. Levy, E. Yahav, [code2vec: learning distributed representations of code](#), Proc. ACM Program. Lang. 3 (POPL)

- (2019) 40:1–40:29. doi:10.1145/3290353.
URL <https://doi.org/10.1145/3290353>
- [30] N. Ahmad, C. Zhang, Enhancing llms interactions for python: A smart api framework for extracting and utilizing semantic code information, in: 2025 5th International Conference on Artificial Intelligence, Big Data and Algorithms (CAIBDA), 2025, pp. 385–388. doi:10.1109/CAIBDA65784.2025.11182836.
 - [31] N. Ahmad, C. Zhang, Enhancing code optimization in llms with dual encoder architecture for syntax and semantic refinement, in: 2025 International Conference on Artificial Intelligence and Digital Ethics (ICAIDE), 2025, pp. 406–409. doi:10.1109/ICAIDE65466.2025.11189637.
 - [32] M. A. Lemley, B. Casey, Fair learning, Texas Law Review 99 (4) (2021) 743–804, available at SSRN: <https://ssrn.com/abstract=3528447>.
URL <https://texaslawreview.org/fair-learning/>
 - [33] N. Vyas, S. M. Kakade, B. Barak, On provable copyright protection for generative models, in: A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, J. Scarlett (Eds.), International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA, Vol. 202 of Proceedings of Machine Learning Research, PMLR, 2023, pp. 35277–35299.
URL <https://proceedings.mlr.press/v202/vyas23b.html>
 - [34] R. D. Cosmo, S. Zacchiroli, Software heritage: Why and how to preserve software source code, in: S. Hara, S. Sugimoto, M. Goto (Eds.), Proceedings of the 14th International Conference on Digital Preservation, iPRES 2017, Kyoto, Japan, September 25-29, 2017, 2017.
URL <https://hdl.handle.net/11353/10.931064>
 - [35] CodeCommons, CodeCommons: Open, responsible, and transparent ai, <https://codecommons.org/>, accessed: 2026-09-03.
 - [36] European Union, Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), official Journal of the European Union, L 1689, 1.8.2024 (August 2024).

URL <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R1689>

- [37] European Commission, Explanatory notice and template for the public summary of training content for general-purpose AI models, <https://digital-strategy.ec.europa.eu/en/library/explanatory-notice-and-template-public-summary-training-content-general-purpose-ai-models>, accessed: 2026-09-03 (Jul. 2025).
- [38] N. Ahmad, C. Zhang, Enhancing code optimization in llms with dual encoder architecture for syntax and semantic refinement, in: 2025 International Conference on Artificial Intelligence and Digital Ethics (ICAIDE), 2025, pp. 406–409. doi:[10.1109/ICAIDE65466.2025.11189637](https://doi.org/10.1109/ICAIDE65466.2025.11189637).