

Towards Code Watermarking with Dual-Channel Transformations

Borui Yang
Shanghai Jiao Tong University
Email: ybirua@sjtu.edu.cn

Wei Li
Shanghai Jiao Tong University
Email: li_wei@sjtu.edu.cn

Liyao Xiang
Shanghai Jiao Tong University
Email: xiangliyao08@sjtu.edu.cn

Bo Li
Hong Kong University of Science and Technology
Email: bli@cse.ust.hk

Abstract—The expansion of the open source community and the rise of large language models have raised ethical and security concerns on the distribution of source code, such as misconduct on copyrighted code, distributions without proper licenses, or misuse of the code for malicious purposes. Hence it is important to track the ownership of source code, in which watermarking is a major technique. Yet, drastically different from natural languages, source code watermarking requires far stricter and more complicated rules to ensure the readability as well as the functionality of the source code. Hence we introduce *SrcMarker*, a watermarking system to unobtrusively encode ID bitstrings into source code, without affecting the usage and semantics of the code. To this end, *SrcMarker* performs transformations on an AST-based intermediate representation that enables unified transformations across different programming languages. The core of the system utilizes learning-based embedding and extraction modules to select rule-based transformations for watermarking. In addition, a novel feature-approximation technique is designed to tackle the inherent non-differentiability of rule selection, thus seamlessly integrating the rule-based transformations and learning-based networks into an interconnected system to enable end-to-end training. Extensive experiments demonstrate the superiority of *SrcMarker* over existing methods in various watermarking requirements.

I. INTRODUCTION

Unauthorized use of source code is an ever-present threat [17], which even worsens with the rapid growth of developer community [18] and recent advances in Large Language Models (LLMs) [42]. For example, a plagiarist copies others’ open-source code, makes minor modifications, and redistributes the code to claim its ownership. Such redistribution requires close inspection as it may well violate the distribution license of the original code. LLMs, on the other hand, can occasionally generate duplicate code snippets [48] from the training dataset which may contain copyrighted code without the original author’s consent. Worse still, for lack of accountability, humans or LLMs could irresponsibly produce plausible code that is faulty or insecure [31], [35], leading to severe consequences if the creator of the code cannot be identified.

Such concerns call for measures to track code provenance, so that the original author could effectively claim its ownership in case of controversies, and platforms could easily identify and filter machine-generated code to prevent misuse [43].

Unfortunately, existing techniques, such as clone detection [7] or authorship attribution [2], [28], are inadequate to serve as direct proofs, as they rely on implicit and passive features (such as code functionality or programming styles [8]) and do not explicitly indicate ownership. Digital watermarking has shown to be a promising solution. It involves a process of hiding a piece of information (i.e., watermark) into digital carriers and later recovering it [47]. The technique has seen wide applications in images [6], [21], texts [11], [58], software [17], [30] and neural networks [3]. Despite recent advances, however, watermarking for source code remains largely unexplored.

Source code distinguishes itself from other content in that it comprises two channels of information: natural and formal [9]. The natural channel is used by developers for comprehension, while the formal channel is utilized by tools such as compilers for automated processing [10]. This intrinsic characteristic determines that code is inherently complicated and delicate. For instance, traditional software watermarking techniques are hindered by the natural channel, as they usually operate on compiled program binaries [15], [17], and hence have no readability requirement. Text watermarking methods, on the other hand, also fall short as they typically rely on word replacement [51] or sentence paraphrase [50], which introduce unrestricted changes that risk breaking the syntax of code.

In this work, we propose *SrcMarker*, a system designed for hiding watermark bitstrings into source code snippets. We embed watermarks at the granularity of functions, so that our system could be scaled to coarser granularities, such as files or projects, by watermarking the functions within them. *SrcMarker* adopts a dual-channel embedding scheme that fits the nature of code: variable name substitutions [59] on the natural channel and semantic-preserving transformations [45] on the formal one. Two channels together enlarge the capacity for watermarking while serving as mutual backups for robustness enhancement. Nonetheless, designing such a system is non-trivial as we must tackle the challenges brought by the dual-channel and discrete nature of code.

The naturalness of code is highly context-dependent, requiring the system to take contextual information into account. Rule-based methods are thus limited as they seldom




Fig. 1. Overview of *SrcMarker*'s embedding and extraction.

consider contexts and tend to produce fixed or homogeneous patterns [54]. To tackle this challenge, we draw inspirations from recent advances in source code processing and design a learning-based backbone for *SrcMarker* to extract contextual features of code. We leverage the rich representation power of neural networks, which have shown to be effective on various source code tasks [20], [38], to capture the contextual information and guide the watermarking process, meanwhile preserving the naturalness of code.

Another challenge lies in the strict grammar rules on the formal channel. Modifications on code must strictly follow its syntax, or otherwise the modified code would be syntactically invalid or semantically incorrect. To this end, we choose to perform semantic-preserving transformations [34] on code for encoding watermark information to maintain operational semantics. However, while such transformations have been widely adopted on source code [10], [54], [55], existing implementations either target limited languages [55] or require external toolchains (such as clang [45] or SrcML [34]) that impact efficiency. Hence we propose *MutableAST*, a unified abstraction of the abstract syntax tree (AST) of the code that allows transformations on different languages to be performed in-memory regardless of language-specific details.

However, both source code and its transformations are discrete and non-differentiable, which are inherently inconsistent with the learning-based design. While existing works on source code have leveraged heuristic [34] or random search [61] to “bypass” such discreteness, these methods could not be used for training a model from scratch as the training process inevitably requires gradient information. To resolve the dilemma, we propose a novel feature-space approximation technique, in which the feature vectors of source code act as a proxy for the discrete operations on code, thus enabling end-to-end training of the entire system.

Figure 1 shows the workflow of *SrcMarker*. The embedding module takes the watermark bitstring and the source code as inputs and outputs the watermarked code with changes in both channels, from which bitstrings are later extracted. Highlights of our contributions are as follows: 1) **Dual-channel source code watermarking**. We propose a deep-learning source code watermarking system conforming to the dual-channel nature of the code, representing one of the first solutions of its kind. 2) **In-memory AST manipulation pipeline**. We introduce *MutableAST*, a language-agnostic pipeline for efficient and reliable

AST manipulation. 3) **End-to-end solution via feature-space approximation**. We propose a novel feature-space approximation technique to harness the discrete nature of source code for end-to-end training. This technique is not only applicable to our design but also to other domains that face similar challenges in source code, such as defending adversarial attacks [45]. Extensive experiments on datasets covering four programming languages (C, C++, Java, JavaScript¹) show that our method could embed and extract watermarks with high accuracy and efficiency, yet incur minimal impact on the natural and operational semantics of code.

II. PROBLEM STATEMENT AND THREAT MODEL

In this section, we first highlight the application scenario, followed by the descriptions on the threat model. We then present our design goals and choices of *SrcMarker*.

Use cases. Watermarking serves as a proactive defense by tracing the provenance of source code at function levels. The code could be a part of, or an entire, open-source project protected by copyright, or a machine-generated snippet by an AI pair programmer. An attacker slightly modifies and redistributes the code without the creator’s consent, violating the license. Or the attacker injects a few lines of faulty, insecure code for malicious spread. In the former case, the owner would collect evidence to claim its ownership of the code and sue the attacker for breach of copyright. In the latter, the owner could act in a responsible manner to detect such misuse of its code and report the event. In both cases, it requires the owner, or an authorized party to identify the ownership from source code.

SrcMarker serves such a purpose by allowing the owner to embed ID bitstrings into its code before the code is released. When encountering suspect code, the owner extracts bitstrings from the suspect using its extraction module.

Threat model. We primarily consider black-box scenarios (e.g., the embedding and extraction modules are guarded by secure black-box APIs), where the adversary has no access to the unwatermarked code, the input watermark, or the model parameters [1] but understands the internal mechanism of *SrcMarker*. The adversary’s goal is to remove the watermark without altering the functionality or impairing the naturalness of the watermarked code, so as to steal and redistribute the watermarked code without being detected. We primarily consider two types of attacks: (1) random changes, where the adversary is aware that a watermark is injected by code transformations, but has no knowledge of the exact transformations and (2) adaptive de-watermarking and re-watermarking attacks, where the adversary has full knowledge on *SrcMarker*’s implementation. More detailed strategies of the adversary will be given in Section IV-E, where we systematically evaluate the robustness of our system.

Design goals. We derive insights from prior works on text watermarking [29] and adapt their requirements to the context of source code, formulating our design goals as follows.

¹In this paper, we use JavaScript and JS interchangeably.

◊ *Effectiveness*. The watermark should be successfully embedded and extracted, preferably with high accuracy and efficiency.

◊ *Transparency*. The watermark should introduce minimal changes and preserve the utility of source code. In this work, we focus on the preservation of operational and natural semantics — if the watermarked code fails to compile or executes differently, it would be useless to anyone.

◊ *Robustness*. Resilience to attacks hinders an adversary from removing or overwriting the watermark and redistributing the code without permission. Typical attacks on source code include semantic-preserving transformations [45] and variable name substitutions [61].

◊ *Capacity*. The watermark, as a unique identifier, should be able to carry sufficiently rich information. The capacity of the watermark is reflected by its length: every additional bit could double the number of representable watermark patterns, yet also demanding a higher capacity of the carrier to incorporate more diversified representations.

Design choices. Taking the design goals and the threat of potential adversaries into consideration, we reveal the choices for our design.

- *Language-agnostic rule-based transformations*. Preservation of utility requires that the modified code must strictly conform to the syntax of the programming language. Modifying the code without concerning any of the lexical, syntactic, and semantic rules would result in compilation failure, let alone keeping its original functionality. Hence we choose rule-based transformations as the major approach to embed watermarks. Such transformations are guided by the structured and formal information from ASTs, thus ensuring the validity and equivalence of the modified code. However, ASTs produced by established parsers include redundant language-specific details that hinder the unified processing of different languages. This would induce additional implementation overhead to fit the language specifics. Therefore, based on ASTs, we introduce *MutableAST*, a new layer of AST abstraction that only retains essential information while hiding the language-specific details, allowing for language-agnostic code manipulation.

- *Dual-channel watermark embedding*. As a salient feature, code conveys information in both formal and natural channels, admitting rule-based modification to both channels. Specifically, we adopt semantic-preserving transformations in the formal channel and variable name substitutions in the natural one. Dual-channel embedding enlarges the possible transformation space, allowing more patterns to be embedded and enhancing watermarking *capacity*. Moreover, the two channels act as mutual backups, enhancing the *robustness* of our system.

- *Learning-based embedding and extraction*. If the embedding and extraction solely rely on rule-based methods, it would require considerable expertise and efforts to design the rules. Moreover, embedding by rules would often generate rigid watermarking patterns, which might negatively impact the *transparency* of watermarks, and less *robust* against adversarial detection. Therefore, we incorporate learning-based embedding and extraction approaches into our system to leverage

the powerful representation capability of neural networks. The networks are combined with rule-based transformations so that the system could provide diversified transformation patterns while still retaining the original meaning in natural and operational channels, thereby ensuring *transparency*.

- *Feature-space approximation for rule-based ops*. Code transformations possess unique characteristics that make their integration with neural networks non-trivial. They operate on discrete source code and cannot be easily modeled as continuous functions. Such a discrete and non-differentiable nature blocks gradient backpropagation and end-to-end training. Since the corresponding feature vectors of source code are continuous, we propose a differentiable proxy for the discrete operations by approximating the differentiable features and thus facilitating end-to-end training.

III. METHODOLOGY

In this section, we first present an overview of *SrcMarker*, and then describe the details of each component.

A. System Overview

Figure 2 illustrates the overall architecture of *SrcMarker*. The system consists of three modules: (1) the watermark embedding module, (2) the watermark extraction module, and (3) the feature approximation module.

The watermark embedding module takes a source code snippet and a watermark bitstring as inputs and outputs a transformed code carrying the watermark information. It embeds the watermark by first using a neural network to select a set of code transformations and variable substitutions, and then performing rule-based transformations accordingly. The watermark extraction module takes the transformed code as input and retrieves the watermark. The feature approximation module serves as a continuous proxy for the non-differentiable code transformations, producing a continuous code feature $\tilde{\mathbf{e}}_{trans}$ that approximates the transformed code feature \mathbf{e}_{trans} , which can be used for watermark decoding. As shown in Figure 2, $\tilde{\mathbf{e}}_{trans}$ provides an alternative differentiable path connecting the code encoder to the watermark decoding loss, which enables the embedding module to update its selections based on the decoding loss.

The training process is guided by three loss functions. First, the actual feature \mathbf{e}_{trans} and the approximated feature $\tilde{\mathbf{e}}_{trans}$ are used to decode the watermarks and compute the actual and approximated decoding loss, $\mathcal{L}_{dec}^{(actual)}$ and $\mathcal{L}_{dec}^{(approx)}$, respectively. To ensure $\tilde{\mathbf{e}}_{trans}$ is a good approximation of \mathbf{e}_{trans} , a third approximation loss \mathcal{L}_{approx} is designed to describe the distance between \mathbf{e}_{trans} and $\tilde{\mathbf{e}}_{trans}$. In summary, *SrcMarker* is trained end-to-end using the combination of the three loss functions,

$$\mathcal{L} = w_1 \mathcal{L}_{dec}^{(actual)} + w_2 \mathcal{L}_{dec}^{(approx)} + w_3 \mathcal{L}_{approx}, \quad (1)$$

where w_1, w_2, w_3 are weight factors for each loss function. More details on the three functions will be given in Section III-C and Section III-D.

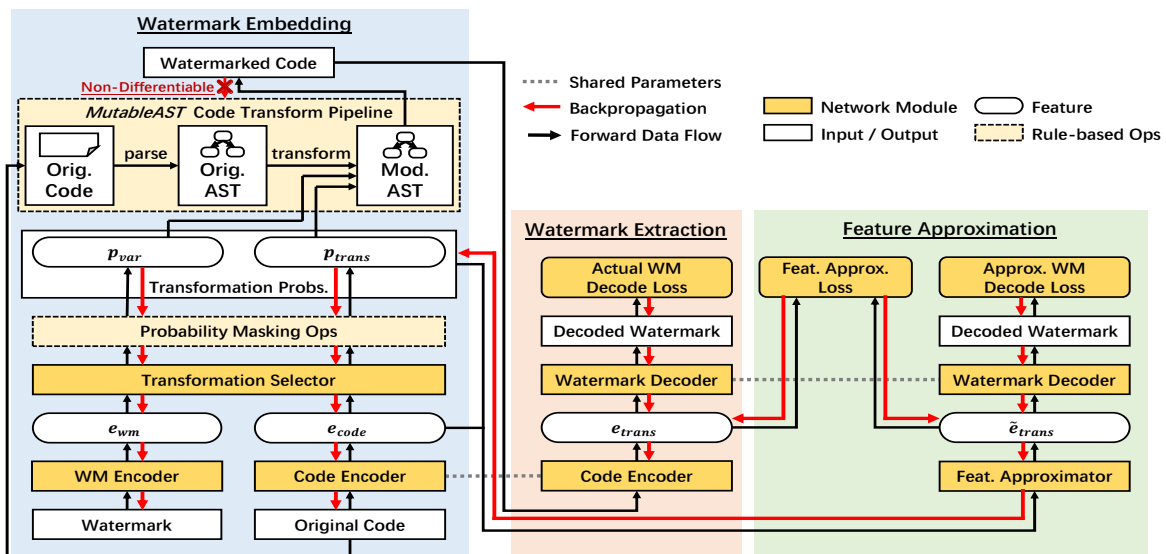


Fig. 2. The architecture of *SrcMarker*. Note that the two “Code Encoder” blocks (in embedding and extraction modules respectively) refer to the same neural network, and the same applies to the two “Watermark Decoder” blocks (in extraction and approximation modules).

At the inference stage, only the watermark embedding and extraction modules are deployed. The transformed code is fed into the extraction module, and only the *actual* transformed feature (e_{trans}) would be used for watermark decoding.

In the following sections, we will first introduce the rule-based code transformation pipeline, dubbed as *MutableAST*, in Section III-B, then describe the watermark embedding and extraction modules in Section III-C, and finally detail the feature approximation module in Section III-D.

B. MutableAST

In this section, we illustrate the intermediate representation (IR) of *SrcMarker* across different programming languages and the semantic-preserving transformations used for preserving the semantics of code.

As aforementioned, we choose ASTs as the IR on which to perform transformations. Such transformations should desirably fulfill two goals: (1) offering a wide variety of options, as the more transformations available, the larger feasible space we have to embed watermarks and (2) being highly efficient in processing, as the code-to-AST, AST modification, AST-to-code conversion would be incorporated iteratively in the training loop as well as the embedding procedure.

Several tools for semantic-preserving transformations have been proposed in prior works. RopGen [34] implements 23 types of transformations using the SrcML² framework, which mainly targets C/C++ and Java; NatGen [10] proposes 6 classes of transformations based on tree-sitter³, which supports 9 different languages.

Unfortunately, existing works have drawbacks that limit their applicability in our scenario. For RopGen, the underlying SrcML framework requires a source code *file* as input and

produces an XML syntax tree, also stored as a *file*. The process involves multiple disk I/Os, which induces significant overhead and is obviously not friendly to the training loop. For NatGen, it uses tree-sitter, which produces language-specific read-only ASTs. Hence one must re-implement the same transformation for every language. Besides, modifications cannot be done directly on the read-only ASTs. Instead, the transformations are performed on token sequences, which becomes cumbersome and error-prone when implementing complex transformations.

To this end, we propose *MutableAST*, a language-agnostic mutable AST abstraction. *MutableAST* is based on tree-sitter and is designed for in-memory parsing. We unify different language-specific ASTs into a language-agnostic representation, by discarding specific details and only retaining information required by code transformations. We store the mutable AST as a special data structure so that modifications can be made in place. *MutableAST* not only allows implementing complex transformations directly on ASTs but also enables unified manipulation on different languages.

The new transformation pipeline is illustrated in Figure 3, where components in gray belong to *MutableAST*. Source code snippets are first parsed into their respective language-specific ASTs by tree-sitter, and then converted to the unified mutable AST by language-specific adaptors. Transformations can then be applied on the mutable AST, and the transformed code can be converted back to source code by language-specific stringifiers, which provide a textual representation of the code with the discarded language details recovered.

Currently, *MutableAST* supports C, C++, Java and JavaScript functions. As listed in Table I, it supports 10 transformation attributes, mostly migrated from or inspired by those in RopGen [34] and NatGen [10]. The transformations range from expression level to block level, which can theoretically provide a maximum of 4096 different combinations of code

²<https://www.srcml.org/>

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

TABLE I
SEMANTIC-PRESERVING TRANSFORMATIONS CURRENTLY SUPPORTED BY *MutableAST*.

Type	Attribute	Description	#Options
Expr	Naming Style	Variable naming style: camelCase, PascalCase, snake_case, _underscore_init	4
Expr	Update Expr.	Styles for increment or decrement expr: i++; ++i, i+=1; i=i+1;	4
Expr	Loop Condition	Conditions for infinite loops: while(true) or while(1)	2
Stmt	Variable Def.	Location of defining local variables: at the beginning of the function or defined on first use	2
Stmt	Variable Init.	Location of initializing local variables: defined and initialized simultaneously or separately: int i=0; or int i; i=0;	2
Stmt	Multiple Defs.	Variables of the same type are defined together or separately: int i, j; or int i; int j;	2
Stmt	Loops	Loop statements: for loops or while loops	2
Stmt	Conditionals	Conditional statements: if structures or switch structures	2
Block	Nested Conditions	Whether multiple if conditions are nested or merged: if(a){if(b)} or if(a && b){}	2
Block	Block Swap	Swapping then and else branches of if statements	2

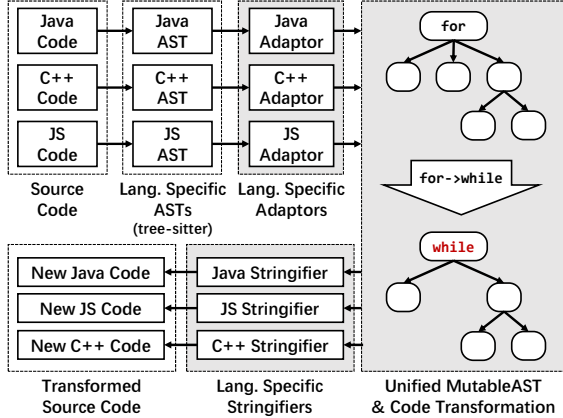


Fig. 3. Overview of the code transformation pipeline in our work. Components in gray belongs to *MutableAST*.

transformations. However, the actual transformation space is still very limited since not all transformations could be applied to a given code snippet (e.g., loop transformations cannot be used on functions with no loops).

C. Watermark Embedding and Extraction

Another major component of *SrcMarker* is learning-based watermark embedding and extraction modules working in both formal and natural channels. The *embedding module* consists of the transformation pipeline described in Section III-B, a source code encoder, a watermark encoder, and a transformation selector.

The *source code encoder* converts the code snippets into feature vectors. We use sequence modeling networks to encode the code snippet. Specifically, the input of the code encoder is a sequence of tokenized source code string $C = [t_1, \dots, t_N]$, where t_i stands for the i -th token. The output is a fixed-length feature vector e_{code} . The *watermark encoder* takes a pre-defined bitstring $B = [b_1, \dots, b_M]$ as the input, and converts the bitstring into the corresponding feature vector e_{wm} . A fully-connected network is used as the watermark encoder. The *transformation selector* takes the code feature e_{code} and the watermark feature e_{wm} as inputs, and outputs two probability vectors, \mathbf{p}_{var} and \mathbf{p}_{trans} , which are the probability over the vocabulary and the probabilistic selection

over all feasible transformation combinations, respectively. The selector contains two independent multi-layer perceptrons (MLPs), f_{var} and f_{trans} , and the probabilities are given by

$$\mathbf{p}_{var} = \text{Softmax}(f_{var}(e_{code} \oplus e_{wm})), \quad (2)$$

$$\mathbf{p}_{trans} = \text{Softmax}(f_{trans}(e_{code} \oplus e_{wm})), \quad (3)$$

where \oplus represents vector concatenation and *Softmax* refers to the softmax activation function.

The *extraction module* consists of a code encoder and a watermark decoder. The module shares exactly the same *code encoder* with the embedding module, which converts watermark code into its feature vector e_{trans} . The *watermark decoder*, which is a fully-connected network, then decodes the watermark using e_{trans} . The loss for watermark decoding is given by the binary cross-entropy loss

$$\mathcal{L}_{dec}^{(actual)} = \text{BCELoss}(f_{wmdec}(e_{trans}), \mathbf{b}) \quad (4)$$

where f_{wmdec} is the watermark decoder and \mathbf{b} is the ground truth watermark bitstring.

D. Feature Approximation

Considered a major contribution of our work, probabilistic modeling and feature approximation are adopted to involve the discrete source code and non-differentiable code transformations in the training loop, enabling end-to-end training of the watermark embedding and extraction modules.

Probabilistic modeling relaxes the discrete selection of code transformations and variable names into continuous probability distributions. Instead of actually selecting the “hard” transformation combination and variable name, the feature selector outputs “soft” probabilistic vectors \mathbf{p}_{var} and \mathbf{p}_{trans} , which are used to sample the actual transformation and variable name. We use Gumbel-Softmax sampling [1], [25] to sample a one-hot vector by the probabilistic vector. By multiplying the one-hot vector with the token set, the corresponding token is selected. Gumbel-Softmax sampling also enables backpropagation through the sampling process with a reparametrization trick [25].

Additionally, inspired by the attention masks in Transformer models [52], we introduce a masking mechanism to eliminate invalid options and encourage diversity. Using variable names as an example, as is shown in Figure 4 (b), we use a validity

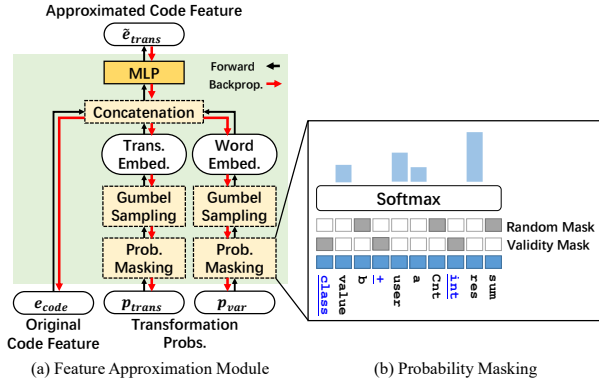


Fig. 4. Network architecture of the feature approximator with probabilistic masking mechanism.

mask to prevent the selection of invalid names, such as operands (+) and keywords (int). Besides, we add another random mask that, similar to the dropout mechanism [23], randomly masks 50% of the variable names, to force the model to explore more options and prevent it from converging to trivial and homogeneous solutions where each watermark bitstring is mapped to a single fixed variable name. The probability for masked logits is 0 after softmax activation, leaving only a proportion of valid names to be selected. A similar validity mask is also applied to the transformation selection to eliminate infeasible transformations. However, we are refrained from applying the additional random mask on transformations as the feasible space is already very limited given a few lines of code.

The *feature approximation module* then bridges the gap between the non-differentiable transformation and the differentiable extraction by aligning \tilde{e}_{trans} with e_{trans} . Intuitively, given a trained encoder, the discrete source code is represented by continuous intermediate representations, and thus the distance in feature vectors corresponds to the dissimilarity between the code. For example, the distance between e_{code} and e_{trans} approximately reveals how much modification was made to the original code. Following this intuition, we implement the feature approximation module, which “predicts” the transformed code feature using the original code feature and the selected transformations. It approximates the effect of code transformations directly in the continuous and differentiable feature space, thus allowing for backpropagation and gradient calculation.

The detailed architecture of the feature approximation module is illustrated in Figure 4 (a), where the aforementioned probability masking is integrated as one of its components. Given the original code feature e_{code} , the probabilities over variable substitution p_{var} and code transformation p_{trans} , it first turns the probabilities into one-hot vectors using Gumbel sampling. The one-hot vectors are then converted to corresponding embedding vectors that represent the selected variable names or transformations. For variable names, we reuse the word vectors from the code encoder; for transfor-

mations, we provide a learnable embedding vector for each transformation combination. The embedding vectors are then concatenated with the original code feature e_{code} , and fed into a fully-connected layer to produce the approximated feature \tilde{e}_{trans} . Note that the entire approximation process is in the differentiable feature space, and thus the gradients can be backpropagated to the embedding module.

To enable backpropagation of the watermark decoding loss to the embedding module, \tilde{e}_{trans} is parallelly used to compute an approximated decoding loss $\mathcal{L}_{dec}^{(approx)}$, given by

$$\mathcal{L}_{dec}^{(actual)} = \text{BCELoss}(f_{wmdec}(\tilde{e}_{trans}), \mathbf{b}). \quad (5)$$

Further, to ensure e_{trans} and \tilde{e}_{trans} are close to each other, we use the mean squared error (MSE) between the approximated and the actual feature as the approximation loss, \mathcal{L}_{approx} , given by

$$\mathcal{L}_{approx} = \|e_{trans} - \tilde{e}_{trans}\|_2^2. \quad (6)$$

IV. EVALUATIONS

In this section, we conduct an empirical evaluation of our system. We first describe the general experimental setup in Section IV-A. Then, we benchmark the efficiency and soundness of *MutableAST* in Section IV-B. For *SrcMarker*, following our design goals, we will evaluate its effectiveness (Section IV-C), transparency (Section IV-D), robustness (Section IV-E) and capacity (Section IV-F). We also evaluate our system in a more realistic project-level watermarking scenario in Section IV-G. Finally, we conduct ablation studies in Section IV-H to validate the components in *SrcMarker*.

A. Experiment Setup

Datasets and pre-processing. Three categories of datasets covering four languages (C, C++, Java and JavaScript) are chosen for evaluation. First, the C++, Java and JavaScript splits of MBXP [4] (MBCPP/MBJP/MBJSP) are used for execution-based evaluation. Originally used for function-level code generation, MBXP contains prompts for code generation queries and corresponding test cases. We repurpose MBXP for watermarking by embedding watermarks into canonical solutions provided by the dataset and verifying the operational semantics of the watermarked code with test cases. Second, the Java and JavaScript splits of CodeSearchNet [24] (CSN-Java/CSN-JS) are used in most of the major experiments on *SrcMarker* due to their extensive uses on large-scale evaluations [26], [38]. They contain Java/JavaScript functions and their corresponding natural language docstrings collected from open-source projects. Additionally, we utilize two smaller datasets, GitHub-C (GH-C) and GitHub-Java (GH-Java), for evaluation on C/C++ languages and experiments at a smaller scale (due to the limitation of baselines). These datasets consist of source files crawled from GitHub, with each file containing one or more functions. Both GH-C and GH-Java have been previously used in RopGen [34]. By default, a 4-bit watermark is embedded into each function.

We list the number of functions, as well as the average number of legal variable names and feasible transformation

TABLE II
DATASETS STATISTICS. LOC STANDS FOR LINES OF CODE.

Dataset	GitHub		MBXP			CSN	
	C	Java	C++	Java	JS	Java	JS
#Functions	4,577	5,501	764	842	797	173,326	63,258
# Valid Vars	5,323	6,361	-	-	-	61,639	85,134
Avg. Transforms	99.5	22.5	61.9	38.0	22.3	18.4	22.5
Avg. LoC.	55.0	45.7	35.9	44.5	30.1	44.7	52.5

combinations on each processed dataset in Table II. The space of variable name substitutes is as large as the token set collected from the training corpus. This statistic for MBXP is missing as we transfer models from other datasets to MBXP and use the token sets of the transferred models. On the other hand, the space of feasible transformation is limited to the effective combinations. On average, transformations on functions from the datasets could theoretically support at least 4-bit watermarks (requiring at least 16 feasible combinations). We further process the datasets by steps and split CSN and GitHub into training/validation/test sets, the detail of which are available in Appendix A-A. Note that MBXP datasets are only used for testing because they are designed for evaluation purposes and are too small to be further divided.

Baselines. For lack of off-the-shelf baselines in source code watermarking, we select two recent natural language watermarking tools, *AWT* [1] and *CALS* [58]. Both methods feature neural networks in their design: *AWT* uses a sequence-to-sequence Transformer for encoding watermarks into texts; *CALS* proposes a context-aware lexical substitution scheme where BERT [16] is used for candidate word generation. However, the original methods do not perform well on source code due to gaps between natural and programming languages. Therefore, we slightly modify the original methods and propose *AWT_{code}* and *CALS_{code}* as baselines. Specifically, *AWT_{code}* shares a similar architecture with *AWT*, but is trained with source code datasets. *CALS_{code}* uses CodeBERT [20] for better adaptation to code. The gaps of the original models and details of our modifications are documented in Appendix A-B.

Implementation details. Our implementation is based on PyTorch. We consider two code encoder variants: a two-layer bidirectional GRU [14] (*SrcMarker_{GRU}*) and a Transformer Encoder [52] with 3 layers and 4 attention heads (*SrcMarker_{TE}*). The feature dimension and embedding size are set to 768. For *SrcMarker_{GRU}*, we use the Adam [32] optimizer with a learning rate 0.001. For *SrcMarker_{TE}*, we use the AdamW [37] optimizer with an exponentially decaying learning rate (decay rate $\gamma = 0.85$) starting from 0.0003, and a weight decay of 0.01. Models are trained on training sets for 25 epochs with a batch size of 64. We use the model with the highest bit accuracy on validation sets for testing and report the performance on test sets. Noted that all experimental results except for Sec. IV-G are reported at the function level.

All experiments are conducted on a Ubuntu 20.04 server with Intel(R) Xeon(R) Gold 6240C CPU @ 2.60GHz with

256GB RAM and an NVIDIA GeForce RTX 3090 GPU.

B. Benchmarking MutableAST

We first evaluate the efficiency and soundness of *MutableAST* as it is responsible for performing transformations in the subsequent watermarking process. The evaluation is performed on MBXP. We compare *MutableAST* with the transformation pipelines proposed in RopGen [34] and NatGen [10], as the two are also featured by semantic-preserving transformations. Note that for RopGen, the results on MBJSP is unavailable because it does not support JavaScript, and that we only consider function-level transformations applicable to C, C++ and Java.

Metrics. The evaluation focuses on the **average runtime per function** and **pass rate** of the transformed code. Pass rate refers to the proportion of transformed code that could pass all unit tests. Since strictly verifying the equivalence of the transformed code typically requires static analysis or formal verification, which are either infeasible on source code snippets or computationally expensive, we leverage the unit tests provided by MBXP to check for correctness. More concretely, for each pipeline, we apply each of the transformations supported by the pipeline on the canonical solutions, and check whether the transformed solutions could pass the test cases.

TABLE III
BENCHMARK OF TRANSFORMATION PIPELINES ON MBXP.

Dataset	Metric	<i>MutableAST</i>	NatGen	RopGen
MBCPP	Pass%	99.93	87.81	99.52
	Time(ms)	0.68	0.77	76.19
MBJP	Pass%	99.95	80.77	98.99
	Time(ms)	0.81	1.06	46.01
MBJSP	Pass%	99.84	87.31	-
	Time(ms)	0.55	0.70	-

Results. The results are reported in Table III. *MutableAST* achieves high efficiency with almost 100% pass rate. Failures are typically due to edge cases of transformations that requires static analysis, which goes beyond the current design of *MutableAST*. In contrast, NatGen, while being comparably efficient, has a much lower pass rate, mainly because it operates on token sequences and is thus more likely to produce errors. RopGen, while also able to achieve high pass rate, is limited by its efficiency. This is mainly due to its frequent file IOs, which would effectively makes the training process more than 50 times slower and is therefore not suitable for our use.

C. Watermark Accuracy and Efficiency

In this section, we evaluate the effectiveness of different watermarking methods.

Metrics. We compare *SrcMarker* with the baselines on accuracy and efficiency. Besides, capacity is also included because unlike *SrcMarker* and *AWT_{code}*, which embed fixed-length watermarks, *CALS_{code}* embed variable-length ones, resulting in differences in capacity. Note that a more comprehensive

TABLE IV

EFFECTIVENESS OF DIFFERENT WATERMARKING METHODS, AVERAGED BY FUNCTIONS. ENCODING TIME (ENC TIME) AND DECODING TIME (DEC TIME) ARE MEASURED IN MILLISECONDS (MS). BPF IS THE NUMBER OF WATERMARK BITS PER FUNCTION.

Dataset	Method	BitAcc	MsgAcc	BPF	EncTime	DecTime
GH-C	<i>SrcMarker_{GRU}</i>	96.19	87.58	4	20.1	7.3
	<i>SrcMarker_{TE}</i>	93.36	79.52	4	12.6	1.2
	<i>AWT_{code}</i>	95.10	81.70	4	36.4	1.8
	<i>CALS_{code}</i>	96.07	92.81	1.22	49680.3	50843.7
GH-Java	<i>SrcMarker_{GRU}</i>	91.97	81.85	4	15.5	6.6
	<i>SrcMarker_{TE}</i>	90.93	75.14	4	9.9	1.4
	<i>AWT_{code}</i>	95.05	82.40	4	37.2	1.7
	<i>CALS_{code}</i>	94.43	91.83	1.40	55722.4	34988.8
CSN-Java	<i>SrcMarker_{GRU}</i>	97.05	91.24	4	14.3	6.4
	<i>SrcMarker_{TE}</i>	97.26	92.74	4	9.4	1.6
	<i>AWT_{code}</i>	92.97	76.30	4	39.0	1.9
CSN-JS	<i>SrcMarker_{GRU}</i>	98.21	94.95	4	18.3	8.3
	<i>SrcMarker_{TE}</i>	96.34	89.84	4	8.3	1.4
	<i>AWT_{code}</i>	89.02	67.07	4	50.1	2.0

evaluation of the capacity of *SrcMarker* will be conducted in Section IV-F.

Accuracy is measured by **bitwise accuracy** (BitAcc) and **message accuracy** (MsgAcc), where BitAcc is the percentage of bits correctly extracted, and MsgAcc is the percentage of watermarks correctly matched. Randomly guessing would result in a BitAcc of 50% and a MsgAcc of 6.25% (for 4-bit watermarks).

Efficiency is measured by the **average time** to embed and extract a watermark. We measure the runtime on the same machine, with each model processing one function at a time. The models will be loaded to GPU in advance, and the time for loading models and data is not included.

Capacity is measured by **bits per function (BPF)**, i.e., the average number of bits embedded into a function.

Results. Table IV reports the average results per function for watermark effectiveness. *SrcMarker* achieves comparable performance with *CALS_{code}* and *AWT_{code}* on the two GitHub datasets, and outperforms the baselines on CSN datasets. The Transformer (*SrcMarker_{TE}*) and GRU (*SrcMarker_{GRU}*) variants have comparable performance. Additionally, *SrcMarker* is more efficient than *AWT_{code}* as the generation process in *AWT_{code}* during watermark embedding would cause slight overhead. *SrcMarker_{TE}* has a shorter embedding and extraction time than *SrcMarker_{GRU}* since the Transformer variant has a slightly smaller model size, and is more suitable for parallel execution. Therefore, *SrcMarker* provides an accurate and efficient solution for source code watermarking.

Despite its accuracy, *CALS_{code}* is limited by its capacity and efficiency. The computation overhead prevents us from running it on CSN-Java (10k samples) or CSN-JS (3.5k samples), as doing so would take approximately 5 days on CSN-JS and even longer for CSN-Java (on a single GPU). Its sequential substitution strategy causes significant overhead: for a function with N tokens, *CALS_{code}* requires $O(N)$ calls to a large language model and it is difficult to parallelize the process due to the sequential order of substitution, thereby

restricting its scalability.

D. Transparency

In this section, we study the transparency of the watermarks by analyzing how much the operational and natural semantics are preserved in the watermarked code.

Metrics. For *operational semantics*, we expect the watermarked code to be equivalent to its original counterpart. However, as aforementioned in Section IV-B, strictly checking semantic equivalence is infeasible on source code snippets. Therefore, we adopt both **syntax check** and **execution-based test** as alternatives. Functions in CSN datasets cannot be executed because they are extracted from projects and do not compile on their own. Therefore, we check whether errors exist in the AST of the watermarked code using tree-sitter. However, the code being syntactically correct is insufficient to guarantee semantic equivalence. To mitigate this, we also evaluate the pass rates by watermarking functions in MBXP and checking whether the watermarked functions could still pass the test cases. Note that MBXP datasets are too small to train models from scratch, so we transfer models trained on other datasets with the same or similar languages: for MBJP and MBJSP, we use models trained on CSN-Java and CSN-JS respectively; for MBCPP, we use models trained on GitHub-C. The transferred models are directly evaluated on MBXP without fine-tuning.

For *natural semantics*, we employ two metrics: **CodeBLEU** [46] and **MRR**. We first directly compare the code before and after watermarking with CodeBLEU, which extends BLEU [44] in NLP by augmenting the n-gram match of BLEU with syntax and dataflow matches. We use CodeBLEU to measure the natural “distance” from the original code and reuse the implementation in CodeXGLUE [38] for computing CodeBLEU. Additionally, A higher level of naturalness means the watermarked code is coherent with the natural language description of its unwatermarked version. Hence we also perform a downstream task — natural language code search on the watermarked code for evaluating naturalness and use the Mean Reciprocal Rank (MRR) as the performance metric. The task searches for the most relevant code in the code corpus on a natural language query [38]. We fine-tune CodeBERT [20] on code search and use the fine-tuned model for computing MRR.

Evaluations on operational semantics are performed on CSN and MBXP, and evaluations on naturalness are conducted on CSN. Note that *CALS_{code}* is not included in the evaluations on CSN because it takes too long to run.

Results for operational semantics. The results for operational semantics are shown in Table V. *AWT_{code}*, despite being trained on dedicated source code datasets, still exhibit critical defects in preserving operational semantics. Few watermarked code could pass the syntax check, not to mention passing the execution-based evaluation. *CALS_{code}* achieves better results on execution-based tests. However, the pass rate is still below satisfactory as more than 25% of the watermarked code fails the test cases. This is because neither method is aware of the

TABLE V

EVALUATION RESULTS ON OPERATIONAL SEMANTICS. FOR CSN, WE USE SYNTAX CHECKS; FOR MBXP, WE USE EXECUTION-BASED CHECKS.

Method		Syntax		Execution		
		CSN-Java	CSN-JS	MBCPP	MBJP	MBJSP
<i>SrcMarker_{TE}</i>	BitAcc	97.26	96.34	96.04	99.44	96.64
	Pass%	100.00	100.00	97.64	97.86	97.99
<i>AWT_{code}</i>	BitAcc	92.97	89.02	97.12	93.88	83.97
	Pass%	0.18	0.51	0.00	0.00	0.00
<i>CALS_{code}</i>	BitAcc	-	-	92.89	93.31	93.50
	Pass%	-	-	68.19	68.65	76.77

intricate rules of programming languages. *AWT_{code}* is only trained to reconstruct the input by heuristically minimizing loss functions, and no explicit constraint is enforced during code generation to ensure grammatical correctness (due to the grammar being too complicated to implement). On the other hand, while it is possible to prevent *CALS_{code}* from substituting keywords with additional rules, the model may still make mistakes such as changing only one occurrence of a variable (which produces an undefined variable). As a result, both models might erroneously break the operational semantics. Failure cases of *AWT_{code}* and *CALS_{code}* are available in Appendix A-C.

In contrast, all code snippets watermarked by *SrcMarker* successfully pass the syntax check, and more than 97% of the watermarked code could pass all test cases on MBXP. Failed samples are mostly corner cases of code transformations or variable substitutions. The AST-based conversion in *SrcMarker* preserves the semantics of the transformed code. Figure 5 demonstrates a snippet before and after watermarking. The variable `pos` in Line 2 is replaced by `circular`, which is a legal variable name, and the `while` loop is changed into a `for` loop equivalently. More importantly, the variable name in Line 17 is changed corresponding to the substitution in Line 3, which shows the power of our intermediate representation. Processing code by sequences, like *AWT_{code}* or *CALS_{code}*, cannot locate the same variables across long distances. Hence, *SrcMarker* encodes the hidden watermark information, meanwhile preserving the operational semantics of the code.

Results for natural semantics. Figure 6 illustrates the results for natural semantics, where higher MRR and CodeBLEU means better preservation of natural semantics. *SrcMarker* achieves a drop of only 0.03 in MRR and maintains a CodeBLEU score above 0.69, which suggests that our method effectively preserves the naturalness of code. On the other hand, *AWT_{code}* exhibit losses in natural semantics. Although *AWT_{code}* achieves a relatively high MRR, its CodeBLEU score is significantly lower due to the presence of syntax errors resulting from watermarking. This is because *AWT_{code}* preserves a large proportion of the tokens and identifiers by which most code searches can be successfully conducted, but the introduced syntax errors would have a negative effect on CodeBLEU since the metric considers syntax in addition to semantic similarity.

```

1 protected void addDelegate(CatalogEntry entry) {
2   int pos = 0;
3   // ...
14  while (local.hasMoreElements()) {
15    if (dp.length() > partial.length()) {
16      pos++;
17    }
18  }
19 }

```

(a) Original code snippet

```

1 protected void addDelegate(CatalogEntry entry) {
2   int circular; // variable renaming
3   circular = 0; // separated decl and init
4   // ...
5   // while-loop -> for-loop
15  for (; local.hasMoreElements(); ) {
16    if (dp.length() > partial.length()) {
17      ++circular; // postfix incr -> prefix incr
18    }
19  }
20 }

```

(b) After watermarking

Fig. 5. A code snippet watermarked by *SrcMarker*. Corresponding changes are highlighted in the same color.

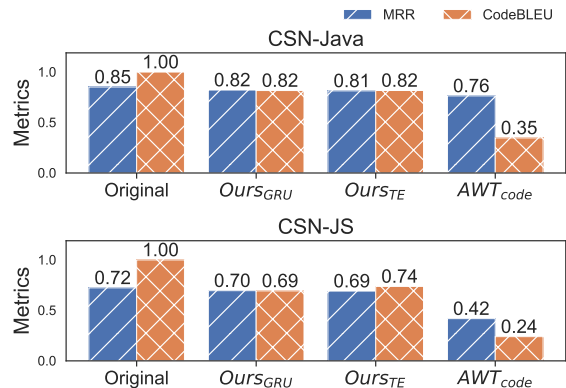


Fig. 6. Natural semantics metrics for *SrcMarker* and *AWT*. “Original” refers to the unwatermarked code.

E. Robustness

In this section, we evaluate how *SrcMarker* performs against an adversary who aims at removing or corrupting the watermark while maintaining the natural and operational semantics of code, so as to redistribute the code without permission. We first consider watermark removal by random changes, and then investigate stronger adaptive attacks.

Random removal attack. We first consider watermark removal attacks, where the adversary is assumed to have general knowledge that watermark is embedded via variable name substitutions and/or code transformations, but has no further knowledge of the exact transformations or variables used. Therefore, the most straightforward attempt to remove the watermark is to randomly perform variable renaming, code transformation, or both. For variable renaming, we consider renaming 25%, 50%, 75%, and 100% of the variables. For code transformations, assume a budget of performing at most 1, 2, or 3 code transformations [34]. For dual-channel attack,

TABLE VI

ROBUSTNESS AGAINST RANDOM REMOVAL ATTACKS. CB STANDS FOR CODEBLEU; “T” STANDS FOR RANDOM CODE TRANSFORMATION; “V” STANDS FOR RANDOM VARIABLE SUBSTITUTION.

Attack	CSN-Java			CSN-JS		
	BitAcc	MRR	CB	BitAcc	MRR	CB
No Atk.	97.26	0.8137	-	96.34	0.6914	-
T@1	83.50	0.8168	0.8589	79.99	0.6970	0.8180
T@2	77.04	0.8159	0.7907	71.54	0.6941	0.7269
T@3	75.07	0.8189	0.7678	67.17	0.6953	0.6900
V@25%	84.08	0.8145	0.9150	85.45	0.6640	0.9063
V@50%	79.89	0.8019	0.8853	82.03	0.6328	0.8693
V@75%	73.71	0.7910	0.8464	74.82	0.6103	0.8238
V@100%	62.68	0.7845	0.7855	64.05	0.5964	0.7538
Dual Ch.	67.76	0.7984	0.7674	67.27	0.6456	0.6951

we consider changing 50% variable names and performing 2 transformations simultaneously.

We apply the above random attacks on watermarked code and report the watermark accuracy after the attack. We also report the MRR and CodeBLEU scores of the watermarked code before and after the attack as measurements of the attack costs.

Results on random removal. Table VI reports results on robustness for $SrcMarker_{TE}$. Results for $SrcMarker_{GRU}$ is similar and is available in Appendix A-D. $SrcMarker$ is relatively robust to random changes on single channels. For moderate attacks such as $\leq 50\%$ variable renaming or ≤ 2 code transformations, the BitAcc is still above 75% for CSN-Java and above 70% for CSN-JS, meaning that most of the watermark information is still retained. As stated in our design choices, dual-channel embedding uses the two channels as mutual backups, thus providing robustness against single-channel attacks. However, if the adversary exploits both channels, or completely removes the information in one channel (e.g., 100% variable renaming), the accuracy would still be negatively impacted. Nonetheless, such attacks comes at a greater cost on natural semantics, resulting in lower MRR or CodeBLEU scores. $SrcMarker_{GRU}$ is actually more robust than $SrcMarker_{TE}$ under random code transformation attacks, possibly attributed to its model structure and training.

Adaptive de-watermarking. We further assume a more knowledgeable adversary who knows about the architecture, training setup, and datasets of $SrcMarker$. We consider adaptive black-box attacks, where the adversary has no access to the model parameters of $SrcMarker$ [1]. Otherwise, we refer to it as a white-box attack. In black-box attacks, the adversary could still leverage its knowledge and train its own model $SrcMarker_{adv}$. In our experiments, we train $SrcMarker_{adv}$ with the same setup as $SrcMarker$ with only differences in initialization. We first consider de-watermarking, where a de-watermarking model takes watermarked code as input, and generates unwatermarked code. Note that in black-box scenarios, the original unwatermarked code is unavailable, so the de-watermarking model is trained on the paired data

(i.e., code before and after watermarking) of $SrcMarker_{adv}$. Both $SrcMarker$ and $SrcMarker_{adv}$ used in this section are the Transformer variants ($SrcMarker_{TE}$). We use a generative GRU network of similar scale as $SrcMarker$ for de-watermarking and evaluate the results on both CSN (for large-scale evaluation) and MBXP (for evaluating the utility of attacked code with execution-based metric).

TABLE VII

ROBUSTNESS AGAINST DE-WATERMARKING. “PASS” IS THE PERCENTAGE OF DE-WATERMARKED CODE THAT PASSES THE TEST CASES. PARENTHESES VALUES REPORT THE ABSOLUTE DROP COMPARED WITH THE WATERMARKED CODE BEFORE ATTACK. BITACC-WB STANDS FOR THE BITACC IN WHITE-BOX SETTINGS.

Metrics	CSN-Java	MBJP	CSN-JS	MBJSP
BitAcc-WB	50.59	50.92	53.27	52.16
BitAcc	73.33 (24%↓)	72.00 (27%↓)	70.13 (26%↓)	72.77 (25%↓)
Pass%	-	5.34 (93%↓)	-	8.03 (91%↓)

Results on de-watermarking. Table VII reports the results on de-watermarking. We first use the de-watermarking model to attack $SrcMarker_{adv}$ (i.e., launching attack in white-box setting) to validate the attack design. This immediately reduces BitAcc to random level (around 50%), indicating the attack is effective. In black-box settings (which our threat model mainly considers), the adversary is able to reduce BitAcc to around 70%. However, most of the de-watermarked code becomes unusable as they cannot pass the original unit tests on MBXP. This is because the de-watermarking model, which uses generative modeling, does not take the syntax rules into consideration (similar to AWT_{code} , the grammar rules are too complex to be integrated into the generation process of the model). We note that for a generative model to produce meaningful and usable code, it should be sufficiently large and trained with adequate amount of data [10], [57]. This is a trade-off made by the adversary. Consider an adversary aiming at removing the watermark from code produced by a large language model. If it had a model large and powerful enough to generate perfectly functioning de-watermarked code, it could directly use its own model to generate code and there would be no need to use the watermarked one [33].

Adaptive re-watermarking. Finally, we consider watermark piracy [19], where the adversary tries to cast ambiguities on the watermarked code by injecting a new watermark using $SrcMarker_{adv}$. We first use $SrcMarker$ to embed a watermark \mathbf{b} , and use $SrcMarker_{adv}$ to embed another watermark \mathbf{b}_{adv} . We use the extraction module of $SrcMarker$ to extract \mathbf{b} and evaluate BitAcc. Note that $SrcMarker_{adv}$ is not used for extraction as it is trained by unauthorized parties and has no credibility.

Results on re-watermarking. Results for re-watermarking are shown in Table VIII. We also first validate the attack under white-box setting (where \mathbf{b}_{adv} is embedded with $SrcMarker$ instead of $SrcMarker_{adv}$). In this case, the BitAcc for \mathbf{b}_{adv} is over 80% while the BitAcc for the original \mathbf{b} is only around 60%, which indicates a successful overwrite and validates the attack. In black-box attacks, the new watermark \mathbf{b}_{adv}

TABLE VIII
EVALUATION RESULTS AGAINST RE-WATERMARKING. WB REFERS TO RESULTS IN WHITE-BOX SETTINGS.

Metrics	CSN-Java	MBJP	CSN-JS	MBJSP
BitAcc-WB (b)	58.95	59.71	60.05	57.72
BitAcc-WB (b_{adv})	83.72	84.47	80.77	82.75
BitAcc (b)	71.14 (26%↓)	71.32 (28%↓)	69.80 (27%↓)	66.50 (31%↓)
BitAcc (b_{adv})	48.74	49.58	43.65	43.10
Pass%	-	97.98% (1%↓)	-	98.49% (1%↓)




Fig. 7. BitAcc changes with respect to watermark capacity (number of watermark bits embedded into a single function).

will also hinder the extraction of **b**. This is because some of the original transformations or variable names used for encoding **b** are overwritten when embedding **b_{adv}**. Further, the re-watermarked code maintains high utility, as **b_{adv}** is also embedded with *MutableAST*. However, the BitAcc for **b_{adv}** in the black-box setting is only at a random level, indicating the overwrite attempt is not successful. An interpretation is that the adversary uses its own model *SrcMarker_{adv}* for embedding **b_{adv}**, which is not trained end-to-end with *SrcMarker*. Consequently, the decoding module of *SrcMarker* will not respond to the changes made by *SrcMarker_{adv}*.

F. Capacity

In this section, we systematically evaluate the capacity of *SrcMarker* at the function level (i.e., the number of bits embedded into a single function) by embedding watermarks of varying lengths (2, 4, 6, or 8 bits) and comparing the BitAcc of *SrcMarker* with *AWT_{code}*. The training setup for the models is the same as in previous sections, except for the number of bits.

Figure 7 illustrates the changes in average BitAcc per function as we vary the bit length. We observe that for all methods, increased capacity comes at the cost of reduced accuracy. Despite this, *SrcMarker* achieves better performance when the number of bits is relatively small (2 or 4 bits) and it still has comparable performance with *AWT_{code}* as the number of bits grows. It is worth mentioning that *SrcMarker* works under tighter constraints as it follows the grammar rules, while *AWT_{code}* could make unrestricted modifications.

G. Project-Level Watermark Verification

Admittedly, the capacity and robustness of watermarks at function level is limited by the few lines of code in a function. To scale the capacity and robustness of our system to more




Fig. 8. Percentage of repositories where the project-level watermark is successfully verified (i.e., null-hypothesis is rejected) under different adversarial settings. “All” refers to the value evaluated on all 213 repositories.

realistic scenarios, we take one step further from the function-level watermarking and apply *SrcMarker* to *project-level watermark verification by function aggregation* [1]. In reality, watermarked functions aggregated at the project level could serve as a more expressive and resilient proof of ownership, as the watermark is sufficiently long as a unique identifier.

The watermark is decomposed into several shorter (e.g., 4-bit) bitstrings, with each embedded into a function of the project. The extracted segments could be concatenated and compared with the original bit sequence. We verify the ID by performing a null hypothesis test according to the number of correctly matched bits [53]. We assume the null hypothesis H_0 to be getting the same number of correctly matched bits (i.e., BitAcc) by chance. Watermark bits generated by randomly guessing would follow a binomial distribution parametrized by $q = 0.5$; for a watermark of length n , the probability of at least k of the bits are correctly matched is

$$\Pr[X > k | H_0] = \sum_{i=k}^n \binom{n}{i} q^i (1-q)^{n-i} = \sum_{i=k}^n \binom{n}{i} 0.5^n. \quad (7)$$

The hypothesis H_0 is rejected if the p -value is below a threshold τ , indicating a low likelihood of getting the same BitAcc by chance. We set $\tau = 0.01$ in our experiments.

Our evaluation is performed on CSN-Java. Each function in CSN-Java belongs to a unique repository. Functions in the test set come from 213 repositories, and we aggregate the watermark segments within each repository to perform verification at the project level. To explore the relationship between verification accuracy and the size of the repository, we further divide the repositories into 4 groups according to the number of functions, namely those containing 1-3, 4-8, 9-33, and more than 34 functions. The group division is determined by the first, second, and third quartile of the number of functions in a single repo of the dataset. We also consider the effect of adversaries by performing random removals as well as adaptive de-watermarking and re-watermarking attacks on the functions.

Figure 8 reports the proportions of repositories where the null hypothesis is rejected, i.e., the project watermark is successfully verified. As the number of watermarked functions grows, the verification success rate also increases. The watermark in a moderately-sized project could be successfully verified with a high likelihood even if the adversary launches strong dual-channel or adaptive attacks. Provided that suffi-

TABLE IX
ACCURACY AND ROBUSTNESS EVALUATION OF *SrcMarker* AND ITS SINGLE-CHANNEL VARIANTS.

Variant	BitAcc	Trans@2	Var@50%
Full	97.26	77.04	79.89
VarOnly	91.25	91.04	68.71
TransOnly	74.22	45.77	72.75

ciently many functions are watermarked, the project watermark could barely be broken without significantly changing every single watermarked function within it, well demonstrating its robustness.

Additionally, to understand the false negative rate of *SrcMarker*, we also perform the test on unwatermarked code, and only 0.94% unwatermarked repositories are verified as watermarked, indicating few unwatermarked repository would be falsely verified. The result is too low to show in Figure 8 and is therefore not plotted.

H. Ablation Study

We analyze some of the design choices made by studying different variants of *SrcMarker*. The default 4-bit embedding setting and GRU code encoder are used.

Dual-channel vs. single channel. We first validate our choice of dual-channel embedding. We train two variants of *SrcMarker* on CSN-Java, which only use the natural channel (variable names) and formal channel (code transformations) to embed watermarks respectively. We compare the BitAcc and robustness against moderate random attacks (namely, renaming 50% variable names, or performing 2 code transformations) of the variants with the original *SrcMarker*.

Table IX reports the results. While merely using variable names could achieve decent accuracy (91.25%), the variant exposes vulnerability to removal attacks that leverage variable renaming: changing 50% of the variable names would reduce its BitAcc to below 70%. On the other hand, only using the formal channel to embed watermarks would result in a BitAcc of only 74.22%, due to the limited space available in transformations. In contrast, by utilizing both channels, *SrcMarker* could achieve high accuracy while maintaining robustness against single-channel attacks.

Random variable mask. As aforementioned, we introduce a random mask to encourage more diversified variable selections. We now validate its effectiveness. We train two *SrcMarker_{GRU}* with and without the random mask respectively and compare their most frequent variable choices for bitstrings 0000 and 1111.

The results are shown in Figure 9. Without random masks, the system would trivially map each bitstring to a single variable name, resulting in repeated patterns that would arguably degrade transparency. Contrary to this, applying random masks could effectively encourage more variable choices, diversifying the patterns of watermarked code.

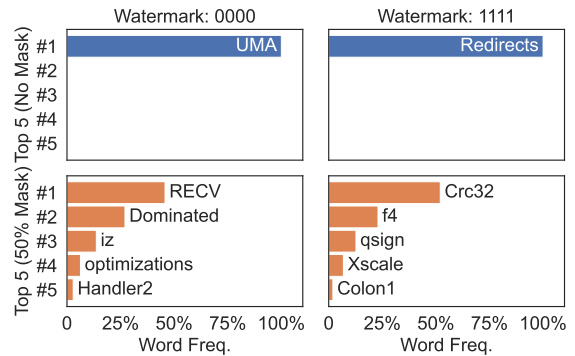


Fig. 9. Top-5 frequently used words for 0000 and 1111 with and without random mask.

V. DISCUSSION

Extendability of *MutableAST*. *MutableAST* should be able to cover common grammar structures for most languages. However, to support a new language, one would have to implement the language-specific adaptors and stringifiers. Additionally, some transformations could be ineffective for certain languages. For example, “Multiple Defs” (in Table I) would have no effect on Python as Python does not explicitly declare variables.

Soundness of *MutableAST*. Despite the fact that the code transformations used in this work are theoretically semantic-preserving [34], we have observed edge cases (in *MutableAST* as well as in NatGen and RopGen) that cause errors. Some are due to bugs in implementation (most of which we have fixed if we were able to identify), but others might require more detailed analysis of the code, which goes beyond the current design of *MutableAST*. Nonetheless, the overall performance of *MutableAST* on the MBXP benchmark is still promising, indicating the current design should be enough for most cases.

Synergy of dual-channel embedding. While we explored the joint utilization of natural and formal channel of code for watermark embedding, our current design have not developed the two channels to their full potential as the information in the two channels are not tightly bound together. Ideally, the two channels could be further entangled together via, for example, feature space regularization [27], to further inter-connect the two channels and bring extra robustness against single-channel attacks.

Verification of operational semantics. We use syntax checks and execution-based tests as substitutions for a thorough check on operational semantics. However, syntax checks alone is far from enough for this purpose. While execution-based metrics tend to be more reliable, there could also be edge cases not covered by test cases. Unfortunately, using formal analysis or software testing would be too expensive.

Transparency of the watermark. In this work, we have primarily focused on the preservation of natural and operational semantics as measures for watermark transparency. Ideally, the watermarked code should be indistinguishable from the original code. However, we have found this goal to be too

ambitious on source code. Even AWT, which has been shown to be stealthy on natural language [1], falls short on code: an automated detector could easily reach a detection F1 score of 0.90 against AWT_{code} on our datasets, which is much higher than the reported F1 score of 0.53 (on natural language) in the original paper. A possible interpretation is that in natural language, there are “redundant” elements (such as articles and prepositions) that are not essential to the meaning of the sentence, and altering such words to encode watermarks would not be evident. Unfortunately, almost every token in source code snippets carry syntactic or semantic meanings, and it is even harder to hide the watermark information due to the strict grammar of programming languages (e.g., when changing a variable name, we must modify every occurrence). Therefore, we leave further improvements on transparency to future work.

VI. RELATED WORK

Software watermarking aims at embedding watermarks into software as a proof of ownership [17]. It can be further divided into static and dynamic watermarking.

Static watermarks are embedded as part of the carrier executable by transforming the file [5], [12], [15]. Kang *et al.* [30] encode watermarks by reordering binary functions. Monden *et al.* [41] embeds watermarks into Java class files by injecting a dummy method to carry the watermark bitstring. However, existing methods are mainly designed for compiled binaries or intermediate representations (e.g., Java bytecode), which do not transfer well to uncompiled source code. Further, most existing works are rule-based, which requires effort and expertise in designing embedding and extraction rules.

Dynamic methods encode watermarks in the runtime behavior of the software. Typically, the watermark can be extracted and verified by triggering a hidden control flow or execution state of the software at runtime [13], [39], [49], [56]. However, such methods require actually executing the software, which is not always possible for uncompiled source code snippets.

Natural language watermarking hides watermarks in texts [29]. Earlier works rely on rule-based modifications, such as synonym substitution [11], [51] and syntactic transformation [40], [50]. More recent works take the powerful natural language processing technique to embed watermarks. Abdelnabi *et al.* [1] propose AWT, a neural text watermarking system based on sequence-to-sequence Transformer [52] models. Yang *et al.* [58] improve substitution-based methods by utilizing large language models [36] for candidate word selection, based on which a sequence incremental watermarking scheme is proposed.

Despite the fact that both source code and text are essentially strings, existing NL watermarking solutions cannot be directly applied to source code. Programming languages have much more sophisticated syntax specifications. Meanwhile, existing methods of text watermarking would apply unrestricted modifications to the input, and might thus break the syntax of the source code, rendering the watermarked code uncompileable.

Semantic-preserving transformations modify the style and structure of code without changing the functionality. Such transformations have seen various applications in deep learning for code. Several works exploit code transformations to craft adversarial examples for source code processing models [45], [59], [61]. Li *et al.* propose RopGen [34], a framework with 23 types of semantic-preserving transformations, for improving the robustness of code authorship attribution models. Further, Chakraborty *et al.* use “recovering unnatural transformations” as a pretraining task and propose NatGen [10], a new large language model for source code. Semantic-preserving transformations offer spaces for modifications while maintaining the validity and equivalence of source code, and thus could represent varied bit strings as watermarks.

The **dual channels of source code** distinguish from natural language: it contains a formal channel and a natural one [9]. The formal channel contains precise, formal semantics, which is typically used by compilers etc. for automatic code processing [10]. The natural channel involves variable names and comments, on which human developers rely to comprehend the code [10], [22]. Our system operates on both channels, which obeys the nature of the code while taking advantage of the ample transformation space in two channels instead of one. We encode watermarks to the two channels in different ways: for the formal one, semantic-preserving transformations are performed under syntactic constraints, while variable name substitution and style modification are adopted for the natural channel.

VII. CONCLUSION

We address the problem of source code ownership verification by proposing *SrcMarker*, a dual-channel source code watermarking system that preserves the validity and semantic equivalence of code. The system is featured by a highly efficient code transformation pipeline integrated into the training loop and a novel feature space approximator to deal with the non-differentiability operations. *SrcMarker* demonstrates superior performance over existing systems in various evaluation aspects. We aim to further improve its secrecy on natural semantics and extend the system to a wider range of programming languages.

REFERENCES

- [1] Sahar Abdelnabi and Mario Fritz. Adversarial watermarking transformer: Towards tracing text provenance with data hiding. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 121–140. IEEE, 2021.
- [2] Mohammed Abuhamad, Tamer AbuHmed, Aziz Mohaisen, and DaeHun Nyang. Large-scale and language-oblivious code authorship identification. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 101–114, 2018.
- [3] Yossi Adi, Carsten Baum, Moustapha Cisse, Benny Pinkas, and Joseph Keshet. Turning your weakness into a strength: Watermarking deep neural networks by backdooring. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*, pages 1615–1631, 2018.

- [4] Ben Athiwaratkun, Sanjay Krishna Gouda, Zijian Wang, Xiaopeng Li, Yuchen Tian, Ming Tan, Wasi Uddin Ahmad, Shiqi Wang, Qing Sun, Mingyue Shang, et al. Multi-lingual evaluation of code generation models. In *The Eleventh International Conference on Learning Representations*, 2022.
- [5] Vivek Balachandran, Ng Wee Keong, and Sabu Emmanuel. Function level control flow obfuscation for software security. In *2014 Eighth International Conference on Complex, Intelligent and Software Intensive Systems*, pages 133–140. IEEE, 2014.
- [6] Shumeet Baluja. Hiding images in plain sight: Deep steganography. *Advances in neural information processing systems*, 30, 2017.
- [7] Lutz Büch and Artur Andrzejak. Learning-based recursive aggregation of abstract syntax trees for code clone detection. In *2019 IEEE 26th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, pages 95–104. IEEE, 2019.
- [8] Aylin Caliskan-Islam, Richard Harang, Andrew Liu, Arvind Narayanan, Clare Voss, Fabian Yamaguchi, and Rachel Greenstadt. De-anonymizing programmers via code stylometry. In *24th {USENIX} Security Symposium ({USENIX} Security 15)*, pages 255–270, 2015.
- [9] Casey Casalnuovo, Earl T Barr, Santanu Kumar Dash, Prem Devanbu, and Emily Morgan. A theory of dual channel constraints. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering: New Ideas and Emerging Results*, pages 25–28, 2020.
- [10] Saikat Chakraborty, Toufique Ahmed, Yangruibo Ding, Premkumar T Devanbu, and Baishakhi Ray. Natgen: generative pre-training by “naturalizing” source code. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 18–30, 2022.
- [11] Ching-Yun Chang and Stephen Clark. Practical linguistic steganography using contextual synonym substitution and a novel vertex coding method. *Computational linguistics*, 40(2):403–448, 2014.
- [12] Jianping Chen, Kui Li, Wanzhi Wen, Weixu Chen, and Chenxue Yan. Software watermarking for java program based on method name encoding. In *Proceedings of the International Conference on Advanced Intelligent Systems and Informatics 2017*, pages 865–874. Springer, 2018.
- [13] Zhe Chen, Chunfu Jia, and Donghui Xu. Hidden path: dynamic software watermarking based on control flow obfuscation. In *2017 IEEE International Conference on Computational Science and Engineering (CSE) and IEEE International Conference on Embedded and Ubiquitous Computing (EUC)*, volume 2, pages 443–450. IEEE, 2017.
- [14] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, 2014.
- [15] Christian Collberg and Tapas Ranjan Sahoo. Software watermarking in the frequency domain: implementation, analysis, and attacks. *Journal of Computer Security*, 13(5):721–755, 2005.
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics, 2019.
- [17] Ayan Dey, Sukriti Bhattacharya, and Nabendu Chaki. Software watermarking: Progress and challenges. *INAE Letters*, 4:65–75, 2019.
- [18] Thomas Dohmke. 100 million developers and counting, 2023. accessed: 2023-06-03. [Online]. Available: <https://github.blog/2023-01-25-100-million-developers-and-counting/>.
- [19] Lixin Fan, Kam Woh Ng, and Chee Seng Chan. Rethinking deep neural network ownership verification: Embedding passports to defeat ambiguity attacks. *Advances in neural information processing systems*, 32, 2019.
- [20] Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, et al. Codebert: A pre-trained model for programming and natural languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1536–1547, 2020.
- [21] Jamie Hayes and George Danezis. Generating steganographic images via adversarial training. *Advances in neural information processing systems*, 30, 2017.
- [22] Abram Hindle, Earl T Barr, Mark Gabel, Zhendong Su, and Premkumar Devanbu. On the naturalness of software. *Communications of the ACM*, 59(5):122–131, 2016.
- [23] Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012.
- [24] Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. Codesearchnet challenge: Evaluating the state of semantic code search. *arXiv preprint arXiv:1909.09436*, 2019.
- [25] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations*, 2016.
- [26] Akshita Jha and Chandan K. Reddy. Codeattack: Code-based adversarial attacks for pre-trained programming language models. In Brian Williams, Yiling Chen, and Jennifer Neville, editors, *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, pages 14892–14900. AAAI Press, 2023.
- [27] Hengrui Jia, Christopher A Choquette-Choo, Varun Chandrasekaran, and Nicolas Papernot. Entangled watermarks as a defense against model extraction. In *USENIX Security Symposium*, pages 1937–1954, 2021.
- [28] Vaibhavi Kalgutkar, Ratinder Kaur, Hugo Gonzalez, Natalia Stakhanova, and Alina Matyukhina. Code authorship attribution: Methods and challenges. *ACM Computing Surveys (CSUR)*, 52(1):1–36, 2019.
- [29] Nurul Shamimi Kamaruddin, Amirrudin Kamsin, Lip Yee Por, and Hameedur Rahman. A review of text watermarking: theory, methods, and applications. *IEEE Access*, 6:8011–8028, 2018.
- [30] Honggoo Kang, Yonghui Kwon, Sangjin Lee, and Hyungjoon Koo. Softmark: Software watermarking via a binary function relocation. In *Annual Computer Security Applications Conference*, pages 169–181, 2021.
- [31] Raphaël Khoury, Anderson R Avila, Jacob Brunelle, and Baba Mamadou Camara. How secure is code generated by chatgpt? *arXiv preprint arXiv:2304.09655*, 2023.
- [32] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [33] John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. A watermark for large language models. *arXiv preprint arXiv:2301.10226*, 2023.
- [34] Zhen Li, Guenevere Chen, Chen Chen, Yayi Zou, and Shouhuai Xu. Ropen: Towards robust code authorship attribution via automatic coding style transformation. In *Proceedings of the 44th International Conference on Software Engineering*, pages 1906–1918, 2022.
- [35] Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation. *arXiv preprint arXiv:2305.01210*, 2023.
- [36] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [37] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018.
- [38] Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin Clement, Dawn Drain, Daxin Jiang, Duyu Tang, et al. Codexglue: A machine learning benchmark dataset for code understanding and generation. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021.
- [39] Haoyu Ma, Chunfu Jia, Shijia Li, Wantong Zheng, and Dinghao Wu. Xmark: dynamic software watermarking using collatz conjecture. *IEEE Transactions on Information Forensics and Security*, 14(11):2859–2874, 2019.
- [40] Hasan M Meral, Emre Sevinc, Bülent Sankur, A Sumru Özsoy, and Tunga Güngör. Syntactic tools for text watermarking. In *Security, Steganography, and Watermarking of Multimedia Contents IX*, volume 6505, pages 339–350. SPIE, 2007.

- [41] A. Monden, H. Iida, K. Matsumoto, K. Inoue, and K. Torii. A practical method for watermarking java programs. In *Proceedings 24th Annual International Computer Software and Applications Conference. COMPSAC2000*, pages 191–197, 2000.
- [42] OpenAI. Introducing chatgpt, 2022. accessed: 2023-06-03. [Online]. Available: <https://openai.com/blog/chatgpt/>.
- [43] OpenAI. Temporary policy: Chatgpt is banned, 2023. accessed: 2023-06-05. [Online]. Available: <https://meta.stackoverflow.com/questions/421831/temporary-policy-chatgpt-is-banned>.
- [44] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002.
- [45] Erwin Quiring, Alwin Maier, Konrad Rieck, et al. Misleading authorship attribution of source code using adversarial learning. In *USENIX Security Symposium*, pages 479–496, 2019.
- [46] Shuo Ren, Daya Guo, Shuai Lu, Long Zhou, Shujie Liu, Duyu Tang, Neel Sundaresan, Ming Zhou, Ambrosio Blanco, and Shuai Ma. Codebleu: a method for automatic evaluation of code synthesis. *arXiv preprint arXiv:2009.10297*, 2020.
- [47] Prabhishik Singh and Ramneet Singh Chadha. A survey of digital watermarking techniques, applications and attacks. *International Journal of Engineering and Innovative Technology (IJEIT)*, 2(9):165–175, 2013.
- [48] Zhensu Sun, Xiaoning Du, Fu Song, Mingze Ni, and Li Li. Coprotector: Protect open-source code against unauthorized training usage with data poisoning. In *Proceedings of the ACM Web Conference 2022*, pages 652–660, 2022.
- [49] Zhenzhou Tian, Qinghua Zheng, Ting Liu, Ming Fan, Eryue Zhuang, and Ziji Yang. Software plagiarism detection with birthmarks based on dynamic key instruction sequences. *IEEE Transactions on Software Engineering*, 41(12):1217–1235, 2015.
- [50] Mercan Topkara, Umut Topkara, and Mikhail J Atallah. Words are not enough: sentence level natural language watermarking. In *Proceedings of the 4th ACM international workshop on Contents protection and security*, pages 37–46, 2006.
- [51] Umut Topkara, Mercan Topkara, and Mikhail J Atallah. The hiding virtues of ambiguity: quantifiably resilient watermarking of natural language text through synonym substitutions. In *Proceedings of the 8th workshop on Multimedia and security*, pages 164–174, 2006.
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [53] Ashish Venugopal, Jakob Uszkoreit, David Talbot, Franz Josef Och, and Juri Ganitkevitch. Watermarking the outputs of structured prediction with an application in statistical machine translation. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1363–1372, 2011.
- [54] Yao Wan, Shijie Zhang, Hongyu Zhang, Yulei Sui, Guandong Xu, Dezhong Yao, Hai Jin, and Lichao Sun. You see what i want you to see: poisoning vulnerabilities in neural code search. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 1233–1245, 2022.
- [55] Deze Wang, Zhouyang Jia, Shanshan Li, Yue Yu, Yun Xiong, Wei Dong, and Xiangke Liao. Bridging pre-trained models and downstream tasks for source code understanding. In *Proceedings of the 44th International Conference on Software Engineering*, pages 287–298, 2022.
- [56] Yilong Wang, Daofu Gong, Bin Lu, Fei Xiang, and Fenlin Liu. Exception handling-based dynamic software watermarking. *IEEE Access*, 6:8882–8889, 2018.
- [57] Yue Wang, Weishi Wang, Shafiq Joty, and Steven CH Hoi. Codet5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8696–8708, 2021.
- [58] Xi Yang, Jie Zhang, Kejiang Chen, Weiming Zhang, Zehua Ma, Feng Wang, and Nenghai Yu. Tracing text provenance via context-aware lexical substitution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11613–11621, 2022.
- [59] Zhou Yang, Jieke Shi, Junda He, and David Lo. Natural attack for pre-trained models of code. In *Proceedings of the 44th International Conference on Software Engineering*, pages 1482–1493, 2022.

TABLE X

NUMBER OF FUNCTIONS IN EACH DATASET DURING PREPROCESSING. *D1* CONTAINS GRAMMATICALLY CORRECT FUNCTIONS. *D2* FURTHER REMOVES FUNCTIONS WITH UNSUPPORTED FEATURES. WE USE *D2* IN OUR EXPERIMENTS.

Dataset	GitHub		MBXP			CSN	
	C	Java	C++	Java	JS	Java	JS
<i>D1</i>	4,731	5,836	766	862	797	180,990	65,200
<i>D2</i>	4,577	5,501	764	842	797	173,326	63,258

- [60] KiYoon Yoo, Wonhyuk Ahn, Jiho Jang, and Nojun Kwak. Robust multi-bit natural language watermarking through invariant features. In Anna Rogers, Jordan L. Boyd-Graber, and Naoaki Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 2092–2115. Association for Computational Linguistics, 2023.
- [61] Huangzhao Zhang, Zhuo Li, Ge Li, Lei Ma, Yang Liu, and Zhi Jin. Generating adversarial examples for holding robustness of source code processing models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 1169–1176, 2020.

APPENDIX A

SUPPLEMENTARY ON EVALUATION

A. Dataset Preprocessing Steps

For typesetting reasons, we leave the details of the steps for processing each of the datasets in this appendix. Since the datasets come in different formats, we first apply specific preprocessing steps to each dataset. Then, we apply a common set of filtering rules to all datasets to obtain the final datasets used in our experiments.

For **MBXP**, since our evaluation requires the canonical solutions, we filter the datasets by removing: (1) samples that do not have a canonical solution and (2) samples whose solution does not pass all test cases (i.e., the solution itself is buggy). For **GitHub-C & GitHub-Java**, since we embed watermarks into functions, but the two GitHub datasets come in the format of files, we extract the functions from each source file. For **CSN-Java & CSN-JS**, since our evaluation on naturalness requires running evaluations on code search, we apply filtering rules from CodeXGLUE [38] to acquire samples used in the code search task.

After the dataset-specific processing, we then use tree-sitter to check and remove functions with grammatical errors. We refer to the filtered datasets as *D1*.

Based on *D1*, we further remove functions containing features not yet supported by *MutableAST*. The final datasets, referred to as *D2*, are used in the experiments in our paper. The statistics of the datasets are shown in Table X. On average, *MutableAST* supports more than 95% of the grammatically correct functions. The rest is left out mostly because it contains features at a coarser granularity than functions but our watermarking is at the function level. For Java, a major unsupported feature is anonymous class declarations; for C/C++, the main unsupported feature is preprocessor directives; for JS, the main unsupported feature is object destructuring patterns.

Finally, we partition the datasets into training, validation and test sets. For GitHub-C and GitHub-Java, we randomly split

them with a ratio of 8:1:1. For CSN datasets, we use their original train/valid/test split (which contain 157,851/4,940/10,535 for CSN-Java and 56,393/3,716/3,149 samples for CSN-JS in D2). However, for MBXP datasets, we use all samples as the test set as they are designed for evaluation purposes only.

B. Implementation Details on Baselines

In this section, we detail the modifications we have made to the baseline models *AWT* and *CALS*.

1) *AWT* and *AWT_{code}*: *AWT* is a sequence-to-sequence Transformer model trained on WikiText-2, a natural language dataset. However, directly transferring the original *AWT* to source code datasets would result in severe out-of-vocabulary (OOV) issues that significantly degrade its performance. Most watermarked snippets would contain the “<unk>” token (a special token representing tokens not included in the model’s vocabulary), which would not only break the syntax of the code but also impair naturalness. For instance, the MRR scores of code watermarked by *AWT* on CSN-Java is only 0.35 (compared with 0.76 of *AWT_{code}* in Figure 6), indicating the natural semantics is barely preserved. To mitigate this, we propose *AWT_{code}*, which has exactly the same architecture as *AWT*, but is trained on the source code datasets to alleviate the OOV issue. *AWT_{code}* uses the same set of hyperparameters as *AWT*, the detail of which could be found in the original paper [1]. To train *AWT_{code}*, we reuse the open-source implementation of *AWT* and slightly modify the data-loading process to port the original implementation to source code datasets.

2) *CALS* and *CALS_{code}*: *CALS* proposes a sequential lexical substitution scheme for encoding watermarks, where a large language model (LLM) is used to generate candidate words for substitution. A synchronicity test will be performed for each generated candidate to ensure the watermark could later be extracted without ambiguity. Additionally, a semantic relatedness (SR) score will be computed by the LLM to ensure the preservation of semantics after substitution. Only candidates who pass the synchronicity test and whose SR scores are above a given threshold could be used for substitution. *CALS* uses BERT [16] for generating candidates and computing SR scores. However, using *CALS* on source code would also meet domain adaptation issues. The candidates are either rejected by the synchronicity test or filtered out due to low SR score, resulting in extremely low watermark capacity. For instance, *CALS* could only embed 0.03 watermark bits per function (BPF) on GitHub-Java, compared with 1.4 BPF for *CALS_{code}* and 4 BPF for *SrcMarker* and *AWT_{code}*. This could be attributed to the fact that BERT is trained on natural language and might not adapt well to source code. Therefore, we replace BERT with CodeBERT [20] for better domain adaptation. We use the implementation released by Yoo *et al.* [60] and change the LLM to CodeBERT in our experiments.

C. Failure Cases of *AWT_{code}* and *CALS_{code}*

Two typical failure cases of *AWT_{code}* and *CALS_{code}* are listed in Figure 10. For *AWT_{code}*, it mistakenly changes

```
// original code snippet
private static void updateCache ( Object key, Object value)

// after watermarking
private static void updateCache ( ) key, ) value )
```

(a) Fail case of *AWT_{code}*

```
// original code snippet
vector<int> sumOfAlternates(vector<int> testTuple) {
    // ...
    sum[0] += testTuple[i];
}

// after watermarking
vector<int> sumOfAlternates(vector<int> testTuple) {
    // ...
    sum[0] += testTuuple[i];
}
```

(b) Fail case of *CALS_{code}*

Fig. 10. Fail cases of *AWT_{code}* and *CALS_{code}* showing the code snippet before (top) and after (bottom) watermarking. Corresponding changes are highlighted in the same color.

a parentheses of the function declarator into a comma, and replaced two type identifiers with parentheses, which breaks the syntax structure of the function signature and results in a compilation error. For *CALS_{code}*, it changes only one occurrence of the variable *testTuple* into *testTuuple*, which is undefined in the context and also causes a compilation error.

D. Robustness Evaluation on *SrcMarker_{GRU}*

In this section, we provide supplementary experimental results for the robustness evaluation on *SrcMarker_{GRU}*. We use the same setup described in Section IV-E.

TABLE XI
PERFORMANCE OF *SrcMarker_{GRU}* UNDER DIFFERENT ATTACK SCENARIOS. CB STANDS FOR CODEBLEU.

Attack	CSN-Java			CSN-JS		
	BitAcc	MRR	CB	BitAcc	MRR	CB
No Atk.	97.05	0.8219	-	98.21	0.6953	-
T@1	85.88	0.8218	0.8588	84.64	0.6853	0.8151
T@2	80.81	0.8208	0.7917	77.39	0.6945	0.7231
T@3	78.77	0.8201	0.7688	74.92	0.6822	0.6858
V@25%	83.28	0.8144	0.9151	86.88	0.6670	0.9041
V@50%	78.51	0.8053	0.8850	82.28	0.6268	0.8669
V@75%	72.51	0.7906	0.8469	75.72	0.6079	0.8210
V@100%	61.23	0.7903	0.7858	65.44	0.5650	0.7500
Dual Ch.	67.81	0.8016	0.7712	65.69	0.6449	0.6886

The results are shown in Table XI. The performance of *SrcMarker_{GRU}* is similar to that of *SrcMarker_{TE}* in terms of robustness. It is relatively robust to single-channel attacks. However, if the attack is performed on both formal and natural channel, our method would suffer more significant accuracy loss.