

PrivCode: When Code Generation Meets Differential Privacy

Zheng Liu, Chen Gong[✉], Terry Yue Zhuo[†], Kecen Li, and, Weichen Yu[‡], Matt Fredrikson[†], Tianhao Wang

University of Virginia

[†]Monash University and CSIRO's Data61

[‡]Carnegie Mellon University

Abstract—Large language models (LLMs) have presented outstanding performance in code generation and completion. However, fine-tuning these models on private datasets can raise privacy and proprietary concerns, such as the leakage of sensitive personal information. Differentially private (DP) code generation provides theoretical guarantees for protecting sensitive code by generating synthetic datasets that preserve statistical properties while reducing privacy leakage concerns. However, DP code generation faces significant challenges due to the strict syntactic dependencies and the privacy-utility trade-off.

We propose PrivCode, the first DP synthesizer specifically designed for code datasets. It incorporates a two-stage framework to improve both privacy and utility. In the first stage, termed “privacy-sanitizing”, PrivCode generates DP-compliant synthetic code by training models using DP-SGD while introducing syntactic information to preserve code structure. The second stage, termed “utility-boosting,” fine-tunes a larger pre-trained LLM on the *synthetic privacy-free* code to mitigate the utility loss caused by DP, enhancing the utility of the generated code. Extensive experiments on four LLMs show that PrivCode generates higher-utility code across various testing tasks under four benchmarks. The experiments also confirm its ability to protect sensitive data under varying privacy budgets. We provide the replication package at the GitHub link.¹

I. INTRODUCTION

The development of large language models (LLMs) has propelled code intelligence into a new era. Advanced open-source code LLMs, such as StarCoder [1], CodeLlama [2], DeepSeekCoder [3], and CodeStral [4], have presented phenomenal performance and even rival human capabilities in tasks like code generation [5], [6], code completion [7], and program-based mathematical reasoning [8]. To better cater to downstream programming scenarios requiring domain-specific expertise, it is common to fine-tune code LLMs on proprietary and sensitive code datasets [9].

[✉] Corresponding Author (Chen Gong). Zheng and Kecen work as independent researchers and remote interns at UVA.

¹<https://github.com/Liuzzyg/PrivCode>

Previous works show that LLMs can memorize content from the training dataset and output it during inference [10], [11], [12]. For example, Carlini et al. [11] show that the GPT-2 language model [13] memorizes and outputs the phone number of an individual named ‘Peter W’ with a crafted prompt. For code LLMs, CodexLeaks [14] found that Codex [15] can reproduce code snippets in verbatim from its training set that contain Personally Identifiable Information (PII). Ziegler et al. [16] found that GitHub Copilot [17] memorizes and reproduces code from its training data, including sensitive information like outdated API keys.

We aim to protect the privacy of code datasets through differentially private (DP) code generation, an approach that generates artificial data, preserving the statistical properties of real data while protecting individual privacy [18], [19], [20], [21]. Specifically, we leverage DP to provide a theoretical guarantee for limiting privacy leakage in the synthesizer’s output. We treat each potential sensitive code snippet as a private individual. The instances of code snippets are presented in Figure 1. Section II-D defines the DP in code generation. To achieve DP, a straightforward approach is to leverage DP-SGD [22] to fine-tune a code LLM. However, several challenges remain to directly adopt this method to code generation.

- *Utility decrease by DP fine-tuning*: DP inevitably reduces model utility because the noise required to satisfy DP guarantees can negatively impact learning processes. Synthetic data from DP fine-tuned models often has lower utility than the original dataset, reducing effectiveness in downstream tasks. Even when using parameter-efficient methods like LoRA [23], DP fine-tuning of LLMs requires training a larger number of parameters than traditional code synthesizers, such as CODEFUSION [24]. For example, a GPT-3 (175B) model with a 144MB LoRA adapter is significantly larger than the 75MB CODEFUSION diffusion model [24]. Consequently, fine-tuning LLMs under a given privacy budget requires more Gaussian noise than fine-tuning traditional code synthesizers [25], [24], resulting in a greater degradation of its utility [18], [26].
- *Strong structural dependencies*: Unlike text, code datasets follow strict syntax, semantic rules, and structural dependencies [27]. Noise injection for DP can disrupt syntax or key identifiers, leading to uncompileable or non-functional code.

DP text synthesis [20], [28] overlooks code-specific positional relationships (e.g., paired “*if-else*” blocks), and gradient noise further hampers syntax learning, weakening the usability of synthetic codes. As we show in Section V-A, directly training DP code synthesizers built on LLMs using DP-SGD [22] can lead to utility bottlenecks and limited adaptability to diverse code structures.

To resolve the aforementioned dilemma, we propose **PrivCode**, the first DP synthesizer specifically tailored for code generation. **PrivCode** introduces a two-stage framework, which breaks the learning process into “*privacy-sanitizing*” and “*utility-boosting*.” The first stage focuses on learning the code characteristics under privacy constraints. As previous works have shown [18], [26], the noise scale is positively correlated with the size of model parameters under a fixed privacy budget. Therefore, we use a junior LLM (one with a smaller size of parameters) to mitigate the negative impact caused by DP noise. In particular, we train synthesizers on sensitive code datasets using DP-SGD at this stage. Additionally, to strengthen the learning of structural dependencies information, **PrivCode** introduces a Privacy-free Syntax-Aware (PrivSA) module. It extracts structural tokens from code snippets, embedding termed adversarial code syntactic information, a supplementary objective designed to counteract the disruption of code structure caused by DP noise, directly into the fine-tuning process, solving the problem of strong structural dependencies of code generation.

The second stage “*utility-boosting*” focuses on boosting the synthetic performance. We know that the post-processing property of DP ensures that operations on DP-compliant outputs do not introduce additional privacy cost [29]. Therefore, we then fine-tune a premium LLM (a powerful, larger-parameter model) on the synthetic code produced after the privacy-sanitizing stage to focus on utility refinement. We note that not all synthetic code snippets from the junior LLM are used to fine-tune the premium LLM. Because the junior LLM’s generative ability is limited under DP, its synthetic codes may be suboptimal, or even include wrong codes. We therefore filter the synthetic code using the following validation and retain only the high-quality portion for fine-tuning the premium LLM. To ensure the functional correctness of synthetic code, we execute the generated code snippets in a controlled environment to filter out those that fail to run, following the execution validation [15], [30]. To ensure the semantic correctness of synthetic code, we summarize each synthetic code snippet into a natural language description. The semantic similarity between this summary and the original prompt is then used to filter out irrelevant code snippets, following the round-trip validation [31], [32]. The filtered high-quality synthetic dataset is used to fine-tune a premium LLM without any DP constraints, thereby avoiding utility degradation. This stage is termed the “*utility-boosting*,” mitigating the synthetic performance degradation.

We evaluate the effectiveness of **PrivCode** by leveraging Qwen2.5-Coder-1.5B [33] as a junior LLM (a smaller-parameter model) in the privacy-sanitizing stage, while

Deepseek-Coder-6.7B-Base [3], Qwen2.5-Coder-7B [33], CodeGemma-7B [34], and CodeQwen1.5-7B [35] as the premium LLMs (a powerful larger-parameter model) in the utility-boosting stage. The utility of synthetic code generated by **PrivCode** are evaluated on well-known benchmarks, including HumanEval [15], MBPP [36], EvalPlus (which includes HumanEval+ and MBPP+) [37], and BigCodeBench [38]. Compared to the baselines, **PrivCode** shows improvements of up to 10.1% in code pass rates for instruction following and up to 14.4% for code completion, across the four benchmarks. Besides, we fine-tune synthesizers on our constructed datasets containing amount of real-world PIIs and injected canary samples. Comparing the leakage rates of canary samples in code generated under a privacy budget of $\epsilon = 4$ with those from the no-DP baseline, **PrivCode** achieves a 0% leakage rate, in contrast to the maximum 100% leakage rate observed in no-DP methods. Ablation studies are performed to emphasize the importance of incorporating PrivSA module into DP fine-tuning and the evolutionary paradigm. In summary, our contributions are three-fold:

- We introduce **PrivCode**, the first DP code generation approach using LLMs, which includes two-stage training, “*privacy-sanitizing*” and “*utility-boosting*” stages.
- **PrivCode** introduces the privacy-free syntax-aware DP fine-tuning on sensitive datasets, incorporating adversarial code syntactic information to enhance the code generation capability under DP.
- Comprehensive evaluations show that **PrivCode** outperforms the baseline across five benchmarks and achieves performance close to methods without any privacy protection mechanisms. **PrivCode** empirically shows excellent ability to protect sensitive code.

II. BACKGROUNDS

This section introduces the brief concepts of LLM-based prompted code generation, abstract syntax trees in code, the privacy leakage concerns in code generation, DP in code generation, and the challenges of DP code generation.

A. LLM-based Prompted Code Generation

LLMs have revolutionized the field of automated code generation, enabling significant advancements in both proprietary and open-source models. Among proprietary models, Claude-4 [41] and GPT-4o [42] represent advanced solutions that synthesize accurate, efficient, and contextually appropriate *code snippet*. As shown in Figure 1, a *code snippet* is one of the most common elements in code generation tasks [15], [36], [38]. It appears as the code solution in instruction following data or as the segment following a given code prompt in completion data.

Currently, decoder-only language model architectures, such as Codex [15], have been shown to outperform encoder-decoder models like CodeT5 [43] on prompted code generation tasks. This class of models follows an autoregressive generation paradigm, predicting the next token based on previously generated tokens. Typically, an autoregressive language

TABLE I: An example of code privacy leakage from the dataset OSS-Instruct PII dataset introduced in Section IV-A. The prompt is a public instruction, while the code snippet contains highlighted private information.

| Source | Text |
|------------|---|
| Train Data | <p>Prompt</p> <p>Write a Python function that takes a list of dictionaries as input, where each dictionary represents a user with various attributes. The function should return a new list of dictionaries, where each dictionary contains only the 'id', 'firstName', 'lastName', and 'email' attributes of the corresponding user. The function should also add a new attribute 'fullName' to each dictionary, which is the concatenation of 'firstName' and 'lastName'.</p> |
| | <pre>\\"\\`python\n`def process_users(users):\n result = []\n users = [\n {\n "id": 57,\n "age": 22,\n "firstName": "Eli",\n "lastName": "Gerry",\n "gender": "female",\n "company": "Assembly",\n "email": "eli.gerry@assembly.com",\n "phone": "+1 (990) 42-211",\n "address": "2 Mihi Place, , New Jersey, 1927",\n }\n]\n\n for user in users:\n user["fullName"] = user["firstName"] + " " + user["lastName"]\n\n result = users\n\n return result\n\n\nprint(process_users(users))\n`\\\"\\`</pre> |

```
Construct an API call in Python that returns information about a given book.

import requests

url = "http://openlibrary.org/api/books"
querystring = {"bibkeys": "ISBN:0201558025", "format": "json", "jscmd": "data"}
headers = { 'cache-control': "no-cache", }
response = requests.request("GET", url, headers=headers, params=querystring)
print(response.text)

from typing import List
def has_close_elements(numbers: List[float], threshold: float) -> bool:
    """ Check if in given list of numbers, are any two numbers closer to
        each other than given threshold.
    """
    >>> has_close_elements([1.0, 2.0, 3.0], 0.5) False
    >>> has_close_elements([1.0, 2.8, 3.0, 4.0, 5.0, 2.0], 0.3) True
    """
    for idx, elem in enumerate(numbers):
        for idx2, elem2 in enumerate(numbers):
            if idx != idx2:
                distance = abs(elem - elem2)
                if distance < threshold:
                    return True
    return False
```

Fig. 1: Examples of code snippets. The above part shows *instruction-following data* [39], while the below part shows *code completion data* [40]. Given the prompt, the highlighted part represents the generated code snippet.

model (e.g., GPT-2 [13]) is trained on an original instruction-code dataset. The model generates code tokens from a prompt using sampling strategies such as Greedy Search [44], Beam Search [45], or Top- k Sampling [46]. Given a public prompt p , the probability distribution of the model’s output code sequence $x = (x_1, x_2, \dots, x_n)$ is:

$$\mathbb{P}(x \mid p) = \prod_{i=1}^n \mathbb{P}(x_i \mid x_1, x_2, \dots, x_{i-1}, p),$$

where $\mathbb{P}(x_i \mid x_1, \dots, x_{i-1}, p)$ represents the probability of generating the i -th token x_i given the previously generated tokens $x_{1:i-1}$ and the prompt p . This paper considers code generation tasks from two common real applications: (1) the instruction-following generation task [39] and (2) the code completion task [40], termed prompted code generation. In particular, the instruction is a natural language task description in the instruction-following generation task, while the prefix includes the function signature and comment-format task description in the code completion task. The instruction and the prefix code header can both be seen as the prompt p , while the output code snippet of each task can be seen as x .

B. Abstract Syntax Trees

Abstract Syntax Tree (AST) is a tree-structured data representation of the syntactic structure of source code snippet [47]. It abstracts the syntactic components of a code snippet while omitting details such as parentheses and whitespace that do not affect the syntax, facilitating code analysis and transformation. In an AST, each node represents a structural element of the code snippet, such as expressions, statements, variables, or functions, forming a hierarchical syntactic representation. ASTs are widely utilized in compilers, interpreters, code analysis tools, code transformation, and vulnerability detection. ASTs enable a structured understanding of code, which helps in identifying complex code patterns and semantic relationships. They also support syntax-aware operations, making them ideal for tasks like automated refactoring or source-to-source translation [48].

By leveraging the AST, the tokens representing the syntactic structures of a code snippet can be automatically extracted. For example, given a code snippet as follows,

Listing 1: Example python code snippet.

```
def is_safe_to_move(dest, loc, closeEnemyLocs):
    moveIn = 1
    for enemy in rg.locs_around(dest,
        filter_out=('invalid')):
        if enemy in closeEnemyLocs:
            if enemy != loc: moveIn = 0
    return moveIn == 1
```

AST constructs a tree representing the syntactic structure of the code. For instance, a node *FunctionDef* points to structural tokens “`def is_safe_to_move(dest, loc, closeEnemyLocs): ...`” that define the hierarchical and semantic relationships within the code. This node resides at the top level of the AST and contains child nodes representing the function’s components, such as its name, arguments, and body. The parser identifies the function definition by recognizing the “`def`” keyword, followed by the function name and parameter list, and groups the subsequent indented statements into the function body. Each of these syntactic elements corresponds to specific nodes in the AST, allowing precise mapping back to the original source code tokens while capturing their structural and semantic roles.

C. Privacy Leakage in Code Generation

Previous studies [18], [19], [20] have presented that fine-tuning models on sensitive training datasets can lead to the leakage of personal private information. Furthermore, prior works [49], [14] present that the code generation model also raises the risk of models memorizing and reproducing sensitive code snippet containing sensitive information in the training dataset. SantaCoder [50] applies PII detection and redaction method only to code snippets within the Stack [51] code dataset to protect private information. It further categorizes these code PIIs into several common types, including *Emails*, *IP addresses*, *Keys*, *Names*, *Usernames*, and *Passwords*.

Following this, we assume that in prompted code generation, the code snippets that constitute the responses are the primary carriers of private information; rather, the prompts do not pose significant privacy leakage risks, referring to the assumption in previous studies [52], [11]. We describe the sensitive code snippet and non-sensitive prompt as follows.

- Prompts are usually descriptive, general-purpose instructions, or prefix headers to guide the model in generating code and are not tied to particular private contexts, making it suitable for public sharing. Prior work [53] from other tasks of LLMs further supports the validity of this hypothesis.
- Code snippets often embed concrete PII, business logic, algorithms, or specific data structures that are part of a company’s intellectual property or private implementation [54]. Automatically generated code could inadvertently expose private information if it reflects sensitive data from the training corpus (e.g., training data or environmental variables).

Hence, this paper focuses on providing DP safeguard for code snippets in code generation contexts while treating the associated prompt as publicly available information.

Table I provides real examples of code privacy leakage. The highlighted portions represent private information, such as *age*, *name*, *gender*, *company*, *email*, *phone number*, and *address*. Table IX presents that the model fine-tuned without any privacy protection directly generates code snippets containing partial private information fragments from the training data. Code snippets generated under DP do not leak private information while still maintaining correctness under privacy budget $\epsilon = \{1, 4, 10\}$. We focus on developing a DP code synthesizer to protect private information in code snippets.

D. Differential Privacy

DP Notion. Differential Privacy (DP) [55] is a privacy-preserving framework that limits how much the presence or absence of a single data point in a dataset can affect a model’s output, thus protecting private information. A randomized algorithm M satisfies (ϵ, δ) -DP if, for any two neighboring datasets D and D' , it holds that:

$$\Pr[M(D) \in O] \leq e^\epsilon \Pr[M(D') \in O] + \delta,$$

where O means any possible output of M . The privacy budget ϵ indicates how much information the algorithm M can reveal; a smaller ϵ implies stronger privacy. The δ can be intuitively

understood as failure probability [55]. Two datasets D and D' are considered neighbors if one can be obtained from the other by removing or replacing a single data point. This paper is the first one to define the DP in code generation. As explained in Section II-C, it focuses on DP code snippets synthesis while treating prompts as public information. We study DP code generation on two code generation tasks: (1) instruction-following code generation and (2) code completion. The notion of neighboring datasets and the protected entity of code snippets differs between these two tasks. We elaborate on the differences as follows.

- **Instruction-following Code Generation:** A data point is defined as a generated code snippet x , derived from a public instruction prompt p . In the DP framework, we define the notion of neighboring datasets D and D' such that they differ by exactly one code snippet x .
- **Code Completion:** A data point is defined as the completed code snippet x , derived from a public prefix code header p . Similarly, by adding or removing one code snippet x from the datasets D , we obtain neighboring datasets D' .

DP provides a theoretical guarantee to measure privacy leakage in the synthesizer’s output, quantifying the risk of revealing private information about real code snippets from synthetic code while still allowing public prompts to facilitate queries and operations.

DP Under Code Correlations. Code contains structural dependencies, such as identifier reuse, logical flow, and imports, that may introduce correlations across tokens. While correlations in text can theoretically weaken DP [56], code snippets are typically much shorter and more modular than long natural-language passages, reducing such effects in practice. Following established practices of applying DP to language-model training [57], we treat each snippet as an indivisible record and rely on DP’s robustness to arbitrary internal structure. Because DP holds under worst-case within-record correlations, snippet-level DP remains a conservative and appropriate privacy formalization for code generation.

DP for Snippet-Level Protection. Our protected unit is a *code snippet*, which aligns with the typical granularity at which models memorize and potentially leak training content. DP is used not for detecting or redacting PII tokens, but to bound the influence of any snippet on the model’s output, limiting memorization-based extraction attacks [11]. The threat model assumes an adversary who queries the trained synthesizer to recover sensitive training snippets. As our goal isn’t to explicitly detect or define PII, we follow prior work [20], while our canary tests employ PII-style tokens, these markers are purely diagnostic and do not define the privacy scope. The DP applies uniformly to all snippets, including proprietary logic, identifiers, and other sensitive patterns.

DP-SGD. In machine learning, DP-SGD [22] incorporates DP into the training of deep learning models, ensuring that the resulting models satisfy formal DP guarantees. Instead of directly using gradients computed from a batch of data samples, DP-SGD first ensures that no single sample overly influences

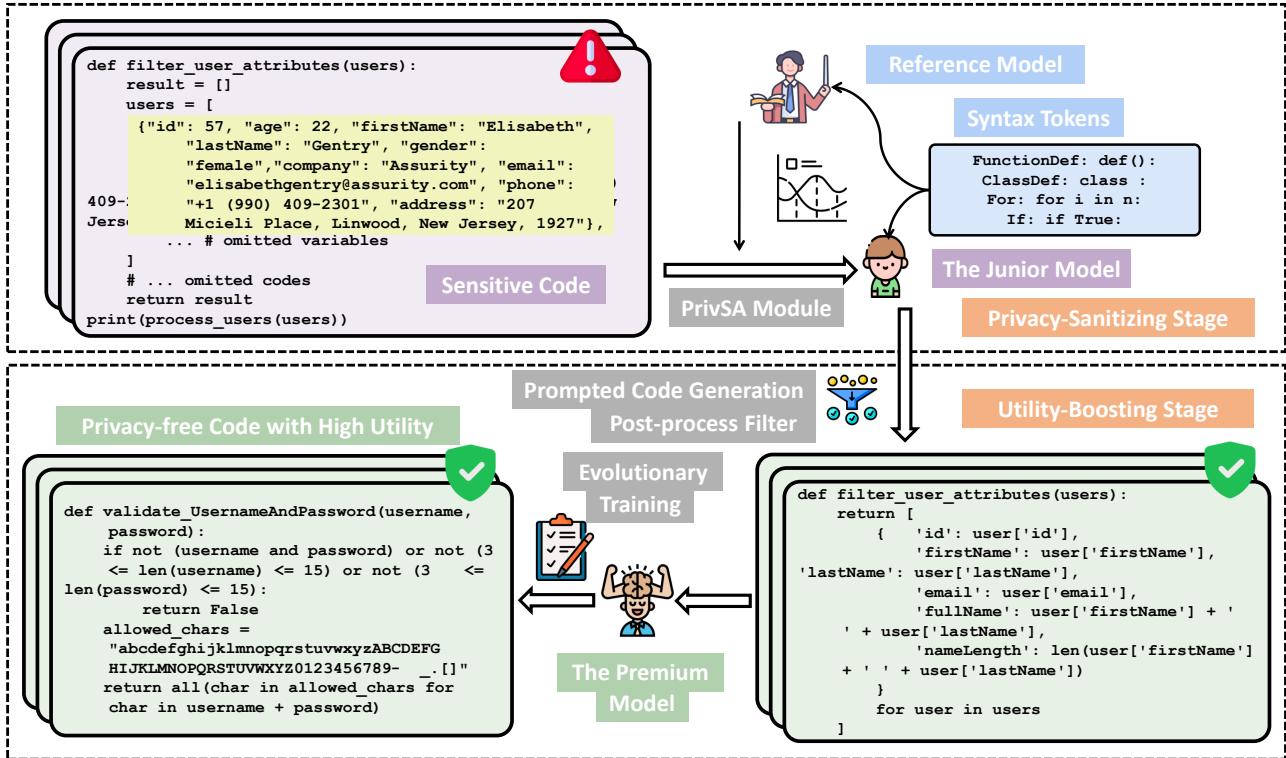


Fig. 2: The workflow of PrivCode. The privacy-sanitizing stage fine-tunes the junior LLM with PrivSA module to generate privacy-free code, incorporating semantic information through knowledge distillation and dynamic adjustment. The utility-boosting stage refines the code via post-process filters to fine-tune the premium LLM and produce high-utility code.

the update step. It achieves this by clipping the ℓ_2 norm of each per-sample gradient g_i to a fixed threshold C : $g_i \leftarrow \text{clip}(g_i, C) = g_i / \max\left(1, \frac{\|g_i\|_2}{C}\right)$. After clipping, DP-SGD adds noise drawn from a Gaussian distribution $\mathcal{N}(0, \sigma^2 \mathbb{I})$ to the averaged clipped gradients: $\tilde{g} = \frac{1}{|B|} \left(\sum_{i \in B} g_i + \mathcal{N}(0, \sigma^2 \mathbb{I}) \right)$, where B is a batch of samples. By controlling both the clipping norm C and the noise scale σ , DP-SGD ensures that the influence of each individual data point on the model parameters remains bounded. A privacy accountant tracks how choices of C , σ , and the number of training steps affect the accumulated privacy loss (ϵ, δ) throughout training, ensuring that the final model respects the desired privacy budget [58].

E. Challenges in DP code generation

This section introduces the challenges in DP code generation from three perspectives.

- **Strong Dependence on Structure.** Code snippets differ from natural language text, as they are structural data governed by strict syntax and semantic rules with well-defined hierarchies. In DP code generation, noise is injected into the code synthesizers during training, potentially disrupting syntax or key identifiers and rendering code un-compilable or compromising its intended functionality.
- **Privacy-Utility Trade-off.** DP mechanisms inherently lower model utility, as the added noise interferes with learning from sensitive data. This leads to synthetic data

from DP fine-tuned models having less utility than the original dataset, reducing effectiveness in downstream tasks.

- **Evaluation and Verification.** Designing a benchmark tailored to real-world code scenarios, embedding private information that LLMs might memorize and reproduce, is essential to judge how effectively DP mechanisms protect sensitive code snippets. The evaluation and verification methods of DP code generation remain largely unexplored.

Code datasets are more structural than text datasets. Injecting noise and applying gradient clipping in DP-SGD disrupt these tokens' positional and semantic associations, impairing the model's specialized understanding and generation of code. In particular, we compute entropy on 5,000 text samples from the Yelp dataset [59] and 5,000 code samples from the code part of Magicoder-OSS-Instruct-75K [60], obtaining values of 6.738 for text and 1.290 for code. Lower entropy indicates more structured data [61].

III. METHODOLOGY

This section introduces PrivCode, the first DP code synthesizer, including adapting existing methods for DP code generation, overview, and technical details.

A. Adapting Existing Methods

A straightforward approach is to adopt DP text synthesis methods, as text and code share similar properties [62], [63]. Previous DP text generation methods, such as AUG-PE [62]

and DP-Prompt [63], progressively guide pre-trained large generative models to produce text synthetic samples resembling sensitive text data. However, these approaches rely solely on pre-trained models without fine-tuning, leading to suboptimal performance in tasks requiring domain-specific expertise. Domain-specific expertise is essential for code generation, as it demands a deep understanding of programming languages, syntax, and functionality. Code outputs must be syntactically correct, semantically meaningful, and executable.

Another approach to DP text generation [64], [20], [65], [66] involves directly training the synthesizer on sensitive text datasets using DP-SGD, making it seemingly adaptable for DP code generation. However, unlike text, code elements such as variables, functions, and control structures are interconnected through strict syntax rules, often reflected in positional relationships (e.g., paired if-else blocks). Traditional methods, without considering such syntax dependencies, make it difficult to learn complex code structures. Besides, the Gaussian noise added to gradients disrupts the learning of fine-grained token-level syntax information and context dependencies, exacerbating errors in the DP synthesizer’s understanding of code syntax and semantics.

B. Overview

This paper proposes PrivCode, addressing the challenges discussed in Section II-E. As the excellent synthesis capabilities of code LLMs across diverse datasets, especially in code generation [2], [35], PrivCode uses code LLMs as the foundational synthesizer. We summarize the high-level contributions of PrivCode as follows.

- **Privacy-free Syntax-aware DP fine-tuning.** As elaborated in Section III-A, code’s structural rules are essential in DP code generation. Besides, the noise introduced by DP-SGD further impedes the synthesizer’s understanding of the code dataset. PrivCode proposes syntax-aware DP fine-tuning, which leverages the syntax information in code datasets to capture the structural correctness intrinsic to code better, thereby enhancing DP code generation. Notably, we extract syntax information from the probability distribution of code tokens, which does not violate DP.
- **Evolutionary paradigm.** Over-relying on one-stage fine-tuning of large-parameter LLMs for DP code generation can lead to excessive utility degradation due to the large Gaussian noise [18], [26]. By leveraging LLMs within a two-stage training framework, defined as the evolutionary paradigm, PrivCode effectively focuses on utility improvement after meeting DP guarantee. In the first stage, we generate synthetic code with DP guarantees. Then, leveraging the post-processing property of DP, we fine-tune a more powerful model on the sanitized outputs without incurring additional privacy cost.

As shown in Figure 2, PrivCode adopts a two-stage approach to separate privacy preservation and generation optimization. The first stage, the “privacy-sanitizing” stage, guides the synthesizer in generating synthetic code under DP. A key property of DP is that operations performed on DP-compliant

Algorithm 1: The workflow of PrivSA module

Input: Input sequence S , junior LLM M_j , reference LLM M_{rf} , upper bound λ_{\max} , lower bound λ_{\min} , decay rate α , step interval Δt
Output: Fine-tuned model M_j^{DP}

```

1 def ASTTokenExtractor ( $S$ ) :
2      $\mathcal{T} \leftarrow \text{ast.parse}(S);$ 
3      $\mathcal{N} \leftarrow \{n \in \mathcal{T} \mid \text{is\_structural}(n)\};$ 
4     return  $\{(\text{type}(n), S[\text{pos}(n)]) \mid n \in \mathcal{N}\};$ 
5 for  $t \in [T]$  do
6     Sample a batch  $L_t$  with probability  $L/N$ ;
    // AST Tokens Extraction
7      $t_s \leftarrow \text{ASTTokenExtractor}(S);$ 
8      $p \leftarrow M_j(t_s);$ 
9      $p' \leftarrow M_{rf}(t_s);$ 
    // Loss computation
10     $L_{\text{KL}} \leftarrow \text{ComputeKLoss}(p, p');$ 
11     $t' = \lfloor t/\Delta t \rfloor \cdot \Delta t$ 
12     $\lambda = \lambda_{\min} + (\lambda_{\max} - \lambda_{\min}) \cdot e^{-\alpha \cdot t'}$ 
    // Gradient Update
13    Update  $M_j$  to minimize Eq. (1) by using DP-SGD;
14 end
15 return  $M_j^{\text{DP}}.$ 

```

outputs do not incur additional privacy loss [29]. This allows synthetic code generated under DP constraints to be further refined without reapplying to privacy mechanisms. In the second stage, termed the “utility-boosting” stage, a larger and more capable LLM is available, as we do not need to consider the privacy here. This premium LLM is fine-tuned on the synthetic code generated by the synthesizer in the “privacy-sanitizing” stage. Then, this stage mitigates the performance degradation caused by DP, enhancing generation quality.

C. Technical Details

Privacy-Sanitizing Stage. This paper introduces the Privacy-free Syntax-Aware (PrivSA) module, which leverages syntax information in code datasets to enhance the DP code synthesizer’s ability to learn code structures. Algorithm 1 presents the processes of PrivSA module. PrivSA module first separately extracts structure tokens from the junior LLM M_j and the reference pre-training LLM M_{rf} . A reference LLM M_{rf} is a parameter-frozen model with the same architecture and weights as M_j . As presented in Line 7, the input is a code snippet sequence, denoted as S . We extract the structural tokens of the code snippet in S using the AST. The sequence S is parsed into an AST, after which structured syntax nodes are filtered out. The type of each node and its corresponding original code tokens are then extracted and returned. This process enables to incorporate structural and syntactic information.

As presented in Lines 8-9, we input the extracted structure tokens t_s into M_j and M_{rf} to obtain the probability distribution p under the current parameters of M_j , as well as the ideal probability distribution p' of M_{rf} . The extent to which p

deviates from p' is calculated by the KL divergence, as shown in Line 10, $\text{KL}(P \parallel Q) = \sum_{x \in \mathcal{X}} P(x) \log \frac{P(x)}{Q(x)}$, where P represents p , and Q represents p' . The KL divergence loss \mathcal{L}_{KL} is scaled by a hyper-parameter λ and combined with the cross-entropy loss \mathcal{L}_{CE} – the standard loss function for language models. The objective of the privacy-sanitizing stage is,

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{CE}} + \lambda \cdot \mathcal{L}_{\text{KL}}. \quad (1)$$

We emphasize that the \mathcal{L}_{KL} is based solely on the probability distribution over extracted structural tokens. And these probability distributions depend only on the internal relationships among widely used public structural tokens, e.g., the “range” token in “`for...in range(n)`”, and do not rely on any relationship with specific code snippets. *It does not involve access to the sensitive training dataset and include additional privacy cost.* Moreover, the gradients of Equation 1, which includes the KL divergence term, are still protected by gradient clipping and noise addition under DP-SGD. Therefore, the PrivSA module does not introduce any additional privacy loss.

As shown in Section V-C, using a constant λ throughout training diminishes the effect of this adversarial information as training iterations increase. The model eventually treats it as a regularization term, leading to results similar to those without incorporating syntactic information. To address this issue, as shown in Line 11-12, we set λ as an exponentially decaying hyper-parameter:

$$\lambda = \lambda_{\min} + (\lambda_{\max} - \lambda_{\min}) \cdot \exp\left(-\alpha \cdot \left\lfloor \frac{t}{\Delta t} \right\rfloor \cdot \Delta t\right), \quad (2)$$

where λ_{\min} and λ_{\max} are the bounds, α controls the decay rate, t is the actual training step, and Δt determines the effective step interval. During the early stages of fine-tuning, leveraging the pre-training model’s capabilities, the structural token position embeddings extracted earlier are better utilized to produce probability distributions closer to the actual tokens, playing a major role in parameter updates. In the later stages of fine-tuning, λ gradually decreases to a predefined lower bound λ_{\min} , stabilizing the training process. Then, the DP code synthesizer updates its parameters using DP-SGD.

Utility-Boosting Stage. Introducing DP into training often results in a decline in the utility of synthetic datasets [18], [67]. To mitigate this, using the post-processing property of DP to refine synthetic data for downstream tasks has become a widely adopted paradigm like previous works [20], [66]. As discussed in Section III-B, PrivCode initially fine-tunes a junior LLM with light parameters on sensitive code, as DP significantly impacts the synthesizer’s learning process. After privacy-sanitizing training, the junior LLM M_j^{DP} is prompted to generate privacy-free code snippets. Leveraging the post-processing property of DP [29], no additional privacy budget is consumed when processing these synthetic code snippets. This dataset is then used to fine-tune a premium LLM M_p with a larger parameter size, enhancing the utility of the synthetic code. PrivCode generates synthetic code using public prompts as conditions, enabling the creation of code

Given a string s , return the length of its longest palindromic subsequence. A subsequence is a sequence that can be derived by deleting some characters without changing the order.

You are given a string `s` consisting of lowercase letters. Your task is to write a function `longestPalindromeSubseq(s)` that finds the length of the longest palindromic subsequence in the given string.

```
def longestPalindromeSubseq(s: str) -> int:
    def helper(b, e):
        if b > e:
            return 0
        if b == e:
            return 1
        if s[b] == s[e]:
            return helper(b + 1, e - 1) + 2
        return max(helper(b + 1, e), helper(b, e - 1))
    return helper(0, len(s) - 1)
```

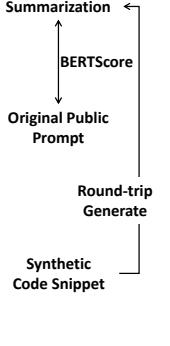


Fig. 3: The workflow of round-trip validation. We prompt a round-trip LLM with synthetic code snippets to generate a natural language summarization. We use BERTScore as a metric to measure the semantic similarity between the summarization and the original prompt.

snippets tailored to specific languages, functionalities, styles, lengths, and structures based on the prompt’s task instructions. This approach results in better correctness and usability code snippets compared to other methods. Furthermore, generating diverse code snippets for the same prompt ensures that the number of generated code snippets meets or exceeds the number of provided public prompts.

However, using the generated code snippets directly for training the premium LLM can result in an unavoidable decline in the performance of the fine-tuned model. This is because some of the snippets are *dirty code snippets*, which may contain incorrect code logic, tokens that do not conform to the programming language’s syntax, extra tokens beyond the complete code, or even mixed natural language tokens.

Execution Validation. Execution validation is a common strategy to ensure that the generated code is syntactically correct and satisfies the target requirements [15], [30]. It involves executing the generated code snippets in a controlled environment to verify their correctness. We perform executable validation by combining heuristic language identification, environment capability probing [68], and sandboxed multi-language runtime execution [69]. Based on runtime signals, compiler diagnostics, dependency checks, and heuristic mismatch detection, we classify failure causes into five major categories as follows.

- Environment error refers to issues caused by incorrect or missing dependencies, invalid API usage, unavailable system resources, or non-existent file paths.
- Compile error indicates that the code fails to compile or initialize the interpreter.
- Runtime error refers to failures occurring after a successful compilation or interpreter startup.
- Language mismatch captures cases where the generated code is inconsistent with specified programming languages.
- Others include empty code generation, timeout failures, or any unspecified abnormal behavior.

Appendix E provides explanations and implementation, including statistical breakdowns of practical failure causes.

Algorithm 2: The Workflow of PrivCode

Input: Sensitive dataset \mathcal{D}_s consists of public prompt and sensitive code pairs $\{(\mathcal{P}, \mathcal{C}_s)\}$, junior LLM M_j , premium LLM M_p , code executor E , round-trip test model M_r , threshold τ_s .

// Fine-Tuning with PrivSA

1 Fine-tune M_j on \mathcal{C}_s using PrivSA module, obtaining the DP code synthesizer M_j^{DP} ;

// Prompted Code Generation

2 $\mathcal{D} \leftarrow \emptyset$;

3 **for** $p \in \mathcal{P}$ **do**

4 | Prompt M_j^{DP} using p to generate privacy-free code snippet c' ; $\mathcal{D} = (p, c') \cup \mathcal{D}$

5 **end**

// Post-Processing Filter

6 **def** ExecutionValidator(\mathcal{D}_i):

7 | $\mathcal{D}_e \leftarrow \emptyset$;

8 | **foreach** $(p_i, c_i) \in \mathcal{D}_i$ **do**

9 | | **if** $E(c_i) = \text{true}$ **then**

10 | | | $c_e \leftarrow c_i$, $d_e \leftarrow (p_i, c_e)$;

11 | | | $\mathcal{D}_e \leftarrow \mathcal{D}_e \cup \{d_e\}$;

12 | | **end**

13 | **end**

14 | **return** \mathcal{D}_e .

15 **def** RoundTripValidator(\mathcal{D}_e):

16 | $\mathcal{D}_f \leftarrow \emptyset$;

17 | **foreach** $(p_i, c_e) \in \mathcal{D}_e$ **do**

18 | | $p_r \leftarrow M_r(c_e)$;

19 | | **if** BERTScore(p_i, p_r) $> \tau_s$ **then**

20 | | | $\mathcal{D}_f \leftarrow \mathcal{D}_f \cup \{(p_i, c_e)\}$;

21 | | **end**

22 | **end**

23 | **return** \mathcal{D}_f ;

24 $\mathcal{D}_e \leftarrow \text{ExecutionValidator}(\mathcal{D})$;

25 $\mathcal{D}_f \leftarrow \text{RoundTripValidator}(\mathcal{D}_e)$;

26 Fine-tune M_p on \mathcal{D}_f without DP;

27 **return** M_p

As presented in Lines 6-14 of Algorithm 2, for each code snippet c_i in the input dataset \mathcal{D}_i , if the execution fails, the snippet is filtered out. The remaining code snippets p_e are combined with their corresponding public prompts p_i as new samples. Only code snippets that pass the execution validation are retained for further processing.

Round-Trip Validation. Even after verifying the correctness of the code itself, there remains a potential issue: the generated code may not match the corresponding instruction well, essentially producing irrelevant or inadequate answers. In other words, the code snippet might fail to satisfy the task requirements of the instruction. To handle this, prior works [31], [70] propose round-trip validation, requiring the model to make predictions (e.g., using natural language to describe some code), provide feedback based on those predictions (e.g., synthesizing code from the predicted description), and check

whether this round-trip process leads to code that is semantically equivalent to the original input, thereby eliminating the need for manual inspection.

Directly comparing generated code with the original sensitive code would introduce additional privacy leakage. We adopt a dual approach by employing a powerful pre-training model M_r to summarize the generated code into natural language descriptions in Line 18. Figure 3 presents the workflow of round-trip validation. The semantic similarity between this summarization and the original prompt is computed using BERTScore [71], which reflects the round-trip correctness of the generated code, i.e., its ability to follow the instructions effectively. In Lines 19- 21, we use a carefully chosen hyper-parameter τ_{sim} as a threshold to filter out samples with low round-trip correctness.

Fine-tuning on Filtered Synthetic Codes. PrivCode fine-tunes the premium LLM M_p on the filtered synthetic code snippets using the regular SGD. The resulting well-tuned LLM is capable of generating high-utility synthetic code while ensuring privacy protection of sensitive code.

IV. EXPERIMENTAL SETUP

This section presents the datasets, baselines, and metrics to assess the utility and privacy protection for DP synthesizers.

A. Utilized Datasets

Our experiments involve two training datasets as follows. We provide more details of studied datasets in Appendix A. **Magicoder-OSS-Instruct-75K.** Magicoder-OSS-Instruct-75K is a dataset created by Magicoder [60] using its OSS-Instruct method for instruction fine-tuning, including a large collection of task-code instruction pairs.

OSS-Instruct PII Dataset. Our constructed instruction-following dataset from the PII dataset [50], containing PII-like private tokens in code snippets.

B. Baselines

To the best of our knowledge, this is the first work to investigate DP code generation. We construct the following three methods as baselines. We provide more details in Appendix B.

- **NonDPFT.** This method directly fine-tunes the premium LLM in PrivCode with LoRA [23], on sensitive code without DP. Without privacy constraints, it is expected to achieve the best utility compared to other methods in downstream tasks. However, the absence of DP increases the risk of privacy leakage from synthetic code.
- **DPFT.** This baseline directly fine-tunes the premium LLM on a sensitive dataset using DP-LoRA [72] under DP. By leveraging DP-SGD [22], DP-LoRA enhance privacy while adding minimal DP noise through fine-tune only parameter-efficient LoRA adapters. However, the added noise inevitably reduces the utility of the synthetic code.
- **DP-Adapter.** DP-Adapter [73] applies DP to standard Adapter tuning [74] by freezing the backbone and updating only the small projection modules. Adapter gradients are optimized via DP-SGD to ensure (ϵ, δ) -DP.

- **JFT.** This method proposes a two-stage fine-tuning framework, termed Just Fine-tune Twice (JFT) [75]. In the first stage, sensitive tokens in the training data are identified using a secret detector and masked, after which the model is fine-tuned on this redacted data without privacy noise. In the second stage, the model is further fine-tuned on the original private data using DP-SGD.

We also list variants of PrivCode that we compare in ablation studies as follows:

- **NonASTPrivCode.** This baseline does not incorporate the code’s syntactic knowledge when fine-tuning the junior model M_j using DP-SGD without PrivSA module.
- **StableASTPrivCode.** This baseline incorporates the code’s syntactic information but sets the KL divergence weight hyper-parameter λ as a constant in PrivSA module, maintaining a fixed adversarial loss for syntactic information. This explores the importance of our carefully designed exponentially decaying λ .
- **NonEvolPrivCode.** This baseline does not use the evolutionary two-stage training framework; instead, it fine-tunes the premium model M_p only in the privacy-sanitizing stage. The aim is to show the importance of evolutionary training.
- **NonPostPrivCode.** This baseline fine-tunes the premium model M_p directly using the synthetic dataset generated by junior model M_j fine-tuned in the privacy-sanitizing stage, demonstrating the necessity of execution validation and round-trip validation for low-quality generated code.

C. Evaluation Benchmarks and Metrics

Utility Evaluation. We select well-known benchmarks such as HumanEval [15], MBPP [36], EvalPlus (including HumanEval+ and MBPP+) [37], and BigCodeBench [38], with their both instruct and complete splits for evaluation. These benchmarks are widely used in evaluating synthetic code [50], [60], [1], [41]. We also select HumanEval-X [76], a multi-programming-language benchmark, to further evaluate the model’s multilingual code generation capabilities. We evaluate our model’s effectiveness by reporting pass rates using greedy decoding, which reduces the impact of randomness, and evaluation results are reported as the pass@1 score metric [15], meaning that the generated code must successfully compile and produce the expected output during execution. The pass@1 score reflects the accuracy of code language models based on their initial attempt to generate the correct code. A higher pass@1 score the synthetic code is better. We also report compile and execution pass rates as utility metrics to support the pass@1 score. The compile pass rate denotes the proportion of generated code snippets that can be successfully parsed and whose target functions can be properly defined, excluding those that fail to compile entirely. The execution pass rate measures the proportion of compiled code snippets that fully pass all tests, calculated relative to the number of successfully compiled samples. We provide the detailed calculated method and more discussions in Appendix C [77].

Private Information Protection Evaluation. We conduct canary experiments to verify whether sensitive training data are

TABLE II: The selection of junior and premium LLM models leveraged in PrivCode. Models are selected based on a balance of code generative ability and parameter size. ‘Base’ means that LLMs are not instruction-tuned.

| Model Type | LLM Model | Year | Type |
|---------------|------------------------------|------|------|
| Junior Model | Qwen2.5-Coder-1.5B [33] | 2024 | Base |
| Premium Model | Deepseek-Coder-6.7B-Base [3] | 2024 | Base |
| | Qwen2.5-Coder-7B [33] | 2024 | Base |
| | CodeGemma-7B [34] | 2024 | Base |
| | CodeQwen1.5-7B [35] | 2024 | Base |

memorized and leaked in the synthetic code snippets generated by PrivCode, following the approach of previous work [20], [10], [78]. We construct five categories of sensitive canary samples and inject them into the OSS-Instruct PII dataset as a training dataset with various repetition rates, each of which contains a distinct type of PII that represents highly private information [79]. The repetition rate refers to the number of times each canary sample is injected into the training dataset. These canary examples, formed as instruction-following pairs, consist of public prompts containing no private information and corresponding code snippet solutions that include one distinct category of PII. Then, using prompts from the testing set of OSS-Instruct PII dataset, we prompt models to generate code snippets. The number of times that the PII canary sample of each category appears in generated code snippets is used to calculate the leakage rate, and category-wise leakage count under different repetition rates. The leakage rate means the distribution of leakage canary categories, while the category-wise leakage count means the number of leakage canary instances of each PII category, both reflecting the ability of methods to protect private information of sensitive datasets.

D. Implementation

Table II shows that PrivCode uses Qwen2.5-Coder-1.5B [33] as the junior model M_j in the privacy-sanitizing stage. We select a smaller parameter model as M_j is to demonstrate that PrivCode’s effectiveness is not dependent on the utility of pre-trained M_j itself. For the round-trip model M_{rt} , we use the powerful Llama-3.1-70B-Instruct [80]. We utilize four commonly used code LLMs as premium models M_p : Deepseek-Coder-6.7B-Base [3], Qwen2.5-Coder-7B [33], CodeGemma-7B [34], and CodeQwen1.5-7B [35]. To accelerate training, in all stages of fine-tuning, we utilize DeepSpeed [81] combined with LoRA [23], a parameter-efficient fine-tuning technique.

V. EMPIRICAL EVALUATIONS

This section first compares the effectiveness of PrivCode with baseline methods in downstream tasks. Next, we analyze how PrivCode protects the sensitive training dataset against privacy leakage. Following that, we examine how hyper-parameters and privacy budgets influence the performance of PrivCode. Finally, we conduct ablation studies to highlight the contributions of the PrivSA module, code post-processing filters, and evolutionary training.

TABLE III: Pass@1 score of PrivCode and baselines under $\epsilon = 4$ trained using four LLMs as premium models. The bolded data represents the best score, and the gray shaded area indicates PrivCode.

| Model | Method | HumanEval | | MBPP | | BigCodeBench | | HumanEval | | MBPP | | BigCodeBench | |
|------------------|-------------|-------------|-------------|-------------|-------------|--------------|-------------|-------------|-------------|-------------|-------------|--------------|------------|
| | | HE | HE+ | MBPP | MBPP+ | Full | Hard | HE | HE+ | MBPP | MBPP+ | Full | Hard |
| | | Instruct | | | | | | | | | | | |
| DS-Coder-6.7B | NonDPFT | 60.4 | 54.9 | 75.1 | 61.4 | 34.5 | 9.5 | 53.0 | 45.7 | 73.8 | 60.1 | 39.5 | 8.1 |
| | DPFT | 54.9 | 48.2 | 57.7 | 48.9 | 28.4 | 4.1 | 45.7 | 39.0 | 65.8 | 54.8 | 34.5 | 8.1 |
| | DP-Adapter | 52.3 | 46.1 | 58.4 | 46.6 | 27.9 | 3.8 | 47.8 | 41.3 | 62.9 | 54.1 | 35.0 | 8.4 |
| | JFT | 53.2 | 48.0 | 60.8 | 52.9 | 28.8 | 5.3 | 49.2 | 42.7 | 66.4 | 52.9 | 31.5 | 8.8 |
| Qwen2.5-Coder-7B | Ours | 56.1 | 51.2 | 69.3 | 59.0 | 29.6 | 7.4 | 47.0 | 43.3 | 69.0 | 58.2 | 36.0 | 9.4 |
| | NonDPFT | 80.5 | 72.6 | 72.5 | 61.9 | 39.0 | 16.2 | 65.2 | 56.7 | 77.2 | 62.7 | 46.1 | 16.2 |
| | DPFT | 59.8 | 53.7 | 58.7 | 51.6 | 22.1 | 5.4 | 36.0 | 29.9 | 56.3 | 46.8 | 26.7 | 6.1 |
| | DP-Adapter | 61.0 | 55.5 | 63.9 | 54.5 | 19.6 | 3.9 | 36.0 | 28.2 | 61.4 | 49.5 | 25.2 | 5.8 |
| | JFT | 63.1 | 56.3 | 69.2 | 60.0 | 21.9 | 7.9 | 38.8 | 33.0 | 67.6 | 54.1 | 27.2 | 8.0 |
| CodeGemma-7B | Ours | 66.5 | 61.0 | 78.3 | 64.8 | 22.9 | 9.5 | 43.9 | 38.7 | 77.9 | 65.6 | 27.9 | 8.8 |
| | NonDPFT | 54.3 | 48.2 | 64.0 | 54.8 | 27.4 | 8.1 | 43.3 | 37.2 | 64.8 | 50.3 | 32.4 | 8.1 |
| | DPFT | 34.1 | 30.5 | 56.1 | 42.9 | 22.6 | 5.1 | 30.5 | 25.6 | 45.0 | 36.2 | 22.7 | 7.4 |
| | DP-Adapter | 36.2 | 29.3 | 58.4 | 44.5 | 20.3 | 4.8 | 34.7 | 25.6 | 48.8 | 40.1 | 24.3 | 7.7 |
| | JFT | 38.5 | 30.8 | 61.7 | 43.6 | 24.9 | 4.8 | 33.2 | 26.9 | 51.7 | 46.6 | 24.8 | 7.2 |
| CodeQwen1.5-7B | Ours | 42.1 | 36.6 | 65.6 | 53.7 | 22.9 | 5.4 | 40.2 | 31.7 | 66.1 | 53.7 | 30.0 | 8.7 |
| | NonDPFT | 64.0 | 56.7 | 73.5 | 61.6 | 33.7 | 10.8 | 59.8 | 53.7 | 74.3 | 60.6 | 39.3 | 12.1 |
| | DPFT | 44.5 | 38.4 | 65.8 | 55.8 | 27.2 | 8.8 | 43.9 | 39.0 | 66.4 | 54.5 | 28.8 | 7.4 |
| | DP-Adapter | 47.1 | 40.2 | 63.7 | 54.9 | 27.6 | 9.5 | 45.2 | 39.8 | 69.6 | 56.1 | 29.8 | 6.7 |
| | JFT | 50.6 | 42.6 | 65.1 | 55.0 | 26.9 | 9.2 | 45.8 | 41.1 | 68.1 | 56.9 | 32.4 | 7.9 |
| Ours | 52.4 | 44.5 | 70.6 | 60.1 | 29.1 | 10.8 | 48.8 | 41.5 | 72.5 | 61.4 | 35.5 | 8.8 | |

TABLE IV: Compile pass rate (Comp.) and execution pass rate (Exec.) of PrivCode and baselines on Qwen2.5-Coder-7B under $\epsilon = 4$ and HumanEval and HumanEval+ benchmarks of instruction-following.

| Method | HumanEval | | | HumanEval+ | | |
|------------|-----------|-------|--------|------------|-------|--------|
| | Comp. | Exec. | Pass@1 | Comp. | Exec. | Pass@1 |
| NonDPFT | 100.0 | 80.5 | 80.5 | 100.0 | 72.6 | 72.6 |
| DPFT | 95.1 | 62.8 | 59.8 | 95.1 | 56.4 | 53.7 |
| DP-Adapter | 97.0 | 61.6 | 61.0 | 95.7 | 58.0 | 55.5 |
| JFT | 98.8 | 64.2 | 63.1 | 97.0 | 57.9 | 56.3 |
| PrivCode | 100.0 | 66.5 | 66.5 | 100.0 | 61.0 | 61.0 |

A. The Utility of Synthetic Codes

Experiment Design. This experiment selects Magicoder-OSS-Instruct-75K dataset as training dataset. We randomly sampled 20,000 examples to train the junior LLM M_j and 55,000 examples for prompted code generation using M_j . After post-processing and validation, the evolutionary training dataset is generated to train the premium LLM M_p . In PrivSA module, we set the upper bound of the KL loss weight hyper-parameter λ_{\max} to 1000, the lower bound λ_{\min} to 0.01, the decay rate α to 0.01, and the step interval Δt to 20. We use slack round-trip parameters, generating 20 round-trip samples per code and setting the similarity threshold τ_{sim} to 0.88 to ensure the selection of high-quality synthetic datasets.

Result Analysis. Table III presents the pass@1 scores of four models fine-tuned with PrivCode and baselines under $\epsilon = 4$. We observe that across most benchmarks, PrivCode consistently achieves the highest pass@1 scores compared to DPFT, DP-Adapter and JFT, and is close to NonDPFT. DPFT and DP-Adapter exhibit similar performance, as both apply parameter-efficient fine-tuning directly leveraging DP-SGD. JFT shows moderate improvement due to more carefully designed training techniques. In the instruct split, among the four models evaluated, PrivCode outperforms the best

TABLE V: Pass@1 score of PrivCode and baselines under $\epsilon = 4$ across Java, C++, and Rust code generation tasks.

| Method | Qwen2.5-Coder-7B | | | CodeGemma-7B | | | CodeQwen1.5-7B | | |
|----------|------------------|-------------|-------------|--------------|-------------|-------------|----------------|-------------|-------------|
| | Java | C++ | Rust | Java | C++ | Rust | Java | C++ | Rust |
| NonDPFT | 55.5 | 17.1 | 50.6 | 42.1 | 17.1 | 31.7 | 54.3 | 22.6 | 39.0 |
| DPFT | 25.0 | 15.9 | 22.0 | 28.0 | 9.1 | 23.8 | 32.3 | 7.9 | 36.0 |
| JFT | 43.9 | 10.4 | 37.7 | 22.6 | 11.6 | 12.8 | 25.0 | 14.0 | 32.3 |
| PrivCode | 57.3 | 22.0 | 44.5 | 42.7 | 12.8 | 28.7 | 54.9 | 17.1 | 42.7 |

baseline with improvements of $\{3.6, 5.8, 9.1, 10.1, 1.9, 2.1\}$ across the HumanEval, HumanEval+, MBPP, MBPP+, Big-CodeBench (full and hard splits) benchmarks. In the complete split, PrivCode achieves improvements of $\{7.0, 5.7, 14.4, 11.5, 5.2, 1.5\}$ over the best baseline on the same set of benchmarks. These results indicate that PrivCode mitigates the adverse impact of DP mechanisms on model utility.

Table IV shows that, across the HumanEval and HumanEval+ benchmarks of instruction-following tasks, PrivCode consistently achieves the highest compile pass rate of 100%, matching the NonDPFT and surpassing all other baselines. Similarly, PrivCode attains the highest execution pass rate compared to baselines, which directly contributes to its superior pass@1 scores. Notably, while other baselines experience a decline in compile pass rate on the more challenging HumanEval+ benchmark—reflecting increased syntax errors, PrivCode maintains zero syntax errors in generated code. This validates the crucial effectiveness of PrivCode in maintaining code syntax correctness under DP.

We also expand experiments to non-Python multi-language code generation tasks, including Java, C++, and Rust code generation tasks as well. Table V further shows that PrivCode consistently achieves the highest multilingual code generation performance across Java, C++, and Rust. Despite the increased syntactic complexity and stricter compiler constraints in these languages, PrivCode outperforms all privacy-preserving base-

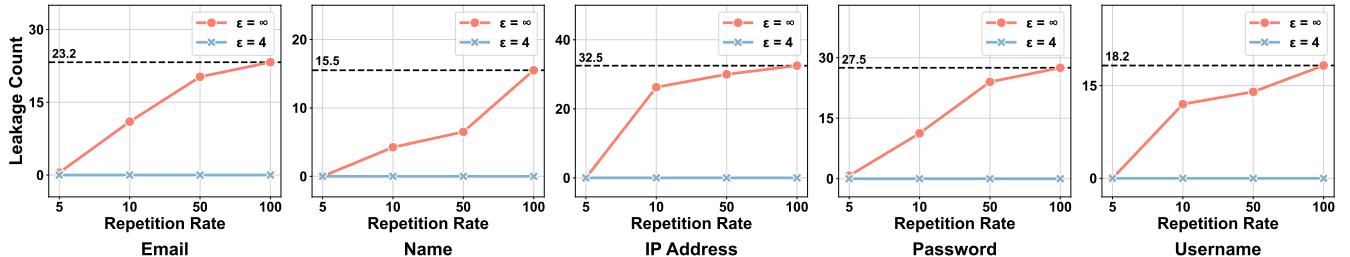


Fig. 4: The average category-wise leakage count of four LLMs trained using PrivCode under $\epsilon = \{4, \infty\}$. The $\epsilon = \infty$ means training PrivCode without DP. Repetition rate is the number of times each canary sample is injected into training datasets.

TABLE VI: The leakage rate of four LLMs trained using PrivCode under $\epsilon = \{4, \infty\}$. Repetition rate refers to the number of times each canary sample is injected into the training dataset.

| Model | Repetition Rate | | | | | |
|------------------|---------------------|----------------|---------------------|----------------|---------------------|----------------|
| | 5 | | 10 | | 100 | |
| | $\epsilon = \infty$ | $\epsilon = 4$ | $\epsilon = \infty$ | $\epsilon = 4$ | $\epsilon = \infty$ | $\epsilon = 4$ |
| DS-Coder-6.7B | 0% | 0% | 40% | 0% | 60% | 0% |
| Qwen2.5-Coder-7B | 20% | 0% | 80% | 0% | 80% | 0% |
| CodeGemma-7B | 0% | 0% | 40% | 0% | 60% | 0% |
| CodeQwen1.5-7B | 40% | 0% | 100% | 0% | 100% | 0% |
| Average | 15% | 0% | 65% | 0% | 75% | 0% |

lines across all settings, achieving maximum improvements of $\{22.6, 6.1, 6.7\}$ pass@1 scores on Java, C++, and Rust respectively, and often approaches the NonDPFT performance. These results confirm that PrivCode effectively preserves the utility of fine-tuning knowledge.

We observe that PrivCode consistently outperforms the DPFT baseline, and JFT also achieves better performance than DPFT on most benchmarks. We analyze that both PrivCode and JFT contain a so-called “utility-boosting stage.” For PrivCode, this refers to fine-tuning on the privacy-free synthetic dataset without DP protection, whereas for JFT, it involves fine-tuning on the masked sensitive dataset, also without DP constraints. These findings highlight the importance of the two-stage paradigm in enhancing model utility under privacy-preserving settings.

The evolutionary training framework presents some intriguing results. For example, Qwen2.5-Coder-7B and CodeQwen1.5-7B achieve pass@1 scores on some benchmarks that even exceed or match NonDPFT (e.g., 65.6 vs. 62.7, 78.3 vs. 72.5, 64.8 vs. 61.9 on MBPP, and 10.8 vs. 10.8 on BigCodeBench). This is because Qwen2.5-Coder-1.5B is used as the junior model M_j , and the premium models M_p , Qwen2.5-Coder-7B and CodeQwen1.5-7B, share similar model architectures and token embeddings in Qwen series. Previous study [82] shows that examples exhibiting similar loss trajectories on small models tend to have similar gradient behavior on larger models, making the synthetic dataset produced by M_j more effective as a high-quality training dataset for fine-tuning M_p to achieve higher utility.

B. Private Information Protection

Experiment Design. The canary experiment is a widely recognized method that inserts unique, private-like sequences,

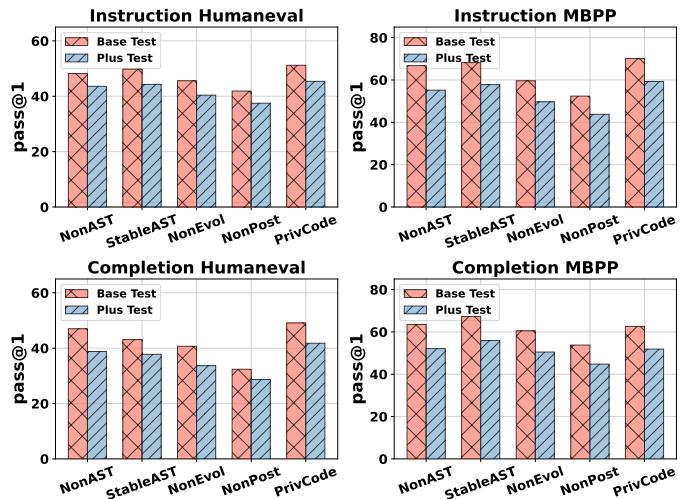


Fig. 5: Average pass@1 scores of four models fine-tuned with PrivCode and its four variants under $\epsilon = 4$. “Base Test” refers to HumanEval or MBPP, while “Plus Test” refers to HumanEval+ and MBPP+.

called canaries, into the training dataset and evaluates whether the model can memorize and reproduce them, assessing the risk of sensitive training information leakage [78], [10], [20].

This experiment uses the OSS-Instruct PII dataset, injected canary samples with varying repetition rates. The evaluation benchmark and metric are described in Section IV-C. We feed Llama-3.1-70B-Instruct [83] with ten instruction-following samples from the OSS-Instruct PII dataset covering different task types as few-shot context to generate five types of canary instruction-following samples. In these canary samples, the instructions are privacy-free task descriptions excluding any PII, while the corresponding code snippets contain the distinct type of PII. The five types of PIIs included are: Email, Name, IP Address, Password, and Username, and each canary sample is restricted to containing only one specific type of PII (e.g., `\\"\\\"\\'python\\n contact = \"sarah.90@gmail.com\\\"\\nprint validate_contact(contact))\\\"\\\"`). We present complete five canary samples and more details in Appendix G in our full version of paper [77].

The hyper-parameter settings are the same as those in Section V-A. To thoroughly explore the variation in PrivCode’s ability to resist privacy data leakage, we conduct this experiment under the privacy budgets $\epsilon = \{4, \infty\}$. To maximize the

TABLE VII: Average Pass@1 of four LLMs trained using **PrivCode**, evaluated on instruct and complete models of HumanEval, MBPP, and EvalPlus benchmarks. The pass@1 score varies with different max lambda. The bolded data is the best score.

| Hyper-Parameter | Model | HumanEval | | MBPP | | HumanEval | | MBPP | |
|----------------------------------|------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | | HE | HE+ | MBPP | MBPP+ | HE | HE+ | MBPP | MBPP+ |
| | | Instruct | | | | Complete | | | |
| $\lambda_{\max} = 1$ | DS-Coder-6.7B | 42.7 | 34.1 | 69.6 | 58.2 | 42.6 | 37.8 | 67.7 | 56.6 |
| | Qwen2.5-Coder-7B | 68.3 | 64.0 | 63.5 | 54.0 | 46.3 | 40.8 | 64.0 | 54.2 |
| | CodeGen-7B | 43.9 | 39.6 | 64.8 | 53.9 | 40.2 | 32.3 | 62.9 | 51.5 |
| | CodeQwen1.5-7B | 43.8 | 40.2 | 67.7 | 58.2 | 49.3 | 42.6 | 68.7 | 59.7 |
| | Average | 49.7 | 44.5 | 66.4 | 56.1 | 44.6 | 38.4 | 65.8 | 55.5 |
| $\lambda_{\max} = 1 \times 10^3$ | DS-Coder-6.7B | 54.3 | 46.3 | 69.9 | 59.8 | 48.2 | 40.2 | 68.0 | 56.9 |
| | Qwen2.5-Coder-7B | 67.1 | 60.4 | 73.0 | 61.9 | 53.7 | 48.2 | 62.2 | 49.7 |
| | CodeGen-7B | 39.6 | 34.8 | 68.0 | 56.9 | 42.1 | 33.5 | 52.1 | 44.4 |
| | CodeQwen1.5-7B | 50.6 | 45.7 | 69.6 | 58.5 | 52.4 | 45.1 | 68.2 | 56.6 |
| | Average | 52.9 | 46.8 | 70.1 | 59.3 | 49.1 | 41.8 | 62.6 | 51.9 |
| $\lambda_{\max} = 1 \times 10^5$ | DS-Coder-6.7B | 35.4 | 33.5 | 69.3 | 58.2 | 46.3 | 39.6 | 69.5 | 57.6 |
| | Qwen2.5-Coder-7B | 73.2 | 64.6 | 76.2 | 63.8 | 46.9 | 42.1 | 73.8 | 63.2 |
| | CodeGen-7B | 42.1 | 37.8 | 62.2 | 48.7 | 35.4 | 27.4 | 57.4 | 46.3 |
| | CodeQwen1.5-7B | 18.9 | 17.7 | 16.1 | 14.0 | 28.0 | 25.0 | 36.5 | 30.1 |
| | Average | 42.4 | 38.4 | 56.0 | 46.2 | 39.2 | 33.5 | 59.3 | 49.3 |

exposure of PII, we set the temperature to 1 and the maximum tokens to 2048 for prompted code generation.

Result Analysis. We start with a case study, as shown in Appendix K of our full version paper [77], which shows that **PrivCode** can avoid leaking private information from the training dataset, while NonDPFT memorizes and reproduces private information from the training data verbatim or partially.

Table VI presents the leakage rate of four LLMs trained using **PrivCode** under $\epsilon = \{4, \infty\}$. We observe that as the repetition rate increases from 5 to 100, the average leakage rate of the four LLMs trained without DP protection rises from 15% to 75%. Notably, CodeQwen1.5-7B exhibits a leakage rate of 100% at repetition rates of 10 and 100, indicating that all types of PII appear in the generated code snippets. This have been demonstrated that repetition in training data is a major contributing factor to model memorization [84]. Importantly, the four LLMs trained with **PrivCode** under $\epsilon = 4$, which represents a relatively strict privacy constraint, show 0% of leakage rate across all repetition rates.

Concretely, Figure 4 shows the average category-wise leakage count of four LLMs trained using **PrivCode** under $\epsilon = \{4, \infty\}$. The average leakage count increases with the repetition rate, and canaries with unique and complex formats, such as Email, IP Address, and Password, are more likely to be memorized by the model. In the case of IP Address, the average leakage count rises from 0 to 26.25, 30 and 32.5 as the repetition rate increases from 5 to 10, 50 and 100. In contrast, under $\epsilon = 4$, the average leakage counts across all five canary categories remain 0, corresponding to a 0% leakage rate. **PrivCode** effectively protects private information from being leaked. Although smaller privacy budgets provide stronger protection, **PrivCode**’s utility inevitably declines as well, as shown in Section V-D. We recommend using a relatively small epsilon, such as 4, which offers sufficient protection for private information while preserving high utility.

C. Abaltion Study

Experiment Design. We adopt stringent configurations of our training pipeline to highlight the importance of **PrivSA**

module, code-specific post-process filters, and the evolutionary training framework. We also demonstrate why the KL divergence weight hyper-parameter λ in **PrivSA** module should be set as described in Section III-C.

Result Analysis. Figure 5 presents the average pass@1 scores of four models fine-tuned with **PrivCode** and its four variants under the instruct and complete modes of four benchmarks.

We observe that **PrivCode** consistently outperforms NonASTPrivCode, with notable improvements on the more complex tasks of HumanEval and HumanEval+, achieving a increase of 3.0 pass@1 scores on both the instruct and complete tests. In Appendix G, we provide a case study: while both **PrivCode** and NonASTPrivCode pass a relatively simple MBPP test case, only **PrivCode** successfully passes a more complex HumanEval test case. StableASTPrivCode, which incorporates constant-proportion adversarial syntactic information, performs better than NonASTPrivCode but gradually degrades into a regularization term during training, ultimately losing effectiveness. In contrast, **PrivCode** achieves up to a improvement of 4.1 in pass@1 scores compared to these syntax-agnostic variants, particularly in the more complex test scenarios of HumanEval+ and MBPP+. These results highlight that **PrivSA** module is the key factor driving **PrivCode**’s superior ability to generate complex code.

Our evolutionary training is the largest contributor to improving code generation utility. As shown, **PrivCode** consistently outperforms NonEvolPrivCode across all benchmarks, with a maximum improvement of 10.5. NonPostPrivCode performs worse than other variants on every benchmark, highlighting the necessity of the utility-boosting stage. Post-processing the privacy-free model’s generated code is essential to ensure a high-quality evolutionary training dataset.

D. Hyper-parameter and Privacy Budget

Experiment Design. We study the impact of two hyper-parameters on synthetic performance: (1) max lambda, $\lambda_{\max} = \{1, 1 \times 10^3, 1 \times 10^5\}$, and (2) privacy budget, $\epsilon = \{0.2, 1, 4, 10, 20\}$. Max lambda represents the proportion of adversarial code syntactic information introduced during the

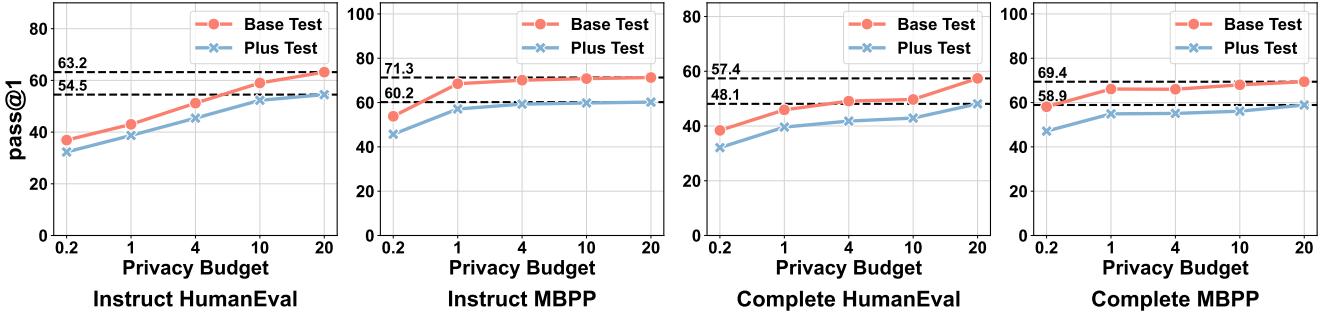


Fig. 6: Average pass@1 of four LLMs trained using PrivCode, evaluated on HumanEval, MBPP, and EvalPlus benchmarks. “Base Test” is HumanEval or MBPP, while “Plus Test” means HumanEval+ and MBPP+. The pass@1 score varies with different privacy budgets. The dashed lines indicate the top value.

privacy-sanitizing stage of PrivCode, especially in the early stages of training the junior model M_j .

Result Analysis. Table VII shows that as the max lambda increases, the average pass@1 scores also rise, reaching their highest on most benchmarks at $\lambda_{\max} = 1 \times 10^3$. However, the scores decline as max lambda continues to increase beyond this point. When max lambda is small, increasing it allows the model to better incorporate the syntactic structure of the code. However, overly large max lambda causes the model to over-focus on local syntactic information, hindering its ability to learn global task-following capabilities and code completeness, leading to overfitting. We find that the optimal parameter settings vary across different trained models. For instance, in the instruct mode of the MBPP benchmark, the optimal λ_{\max} for DS-Coder-6.7B, CodeGemma-7B, and CodeQwen1.5-7B is 1×10^3 , while for Qwen2.5-Coder-7B, it is 1. Differences in model architecture and parameter size lead to varying cross-entropy loss and adversarial KL loss proportions, making a suitable max lambda crucial for each model and dataset.

Figure 6 illustrates that across all benchmarks, the average pass@1 score decreases as the privacy budget decreases, indicating the utility of synthetic code snippets.

VI. RELATED WORK

We discuss related work of PrivCode: code generation (no-DP) and DP dataset synthesis across various fields.

Code Generation. General-purpose LLMs often struggle with code generation due to their emphasis on natural language, leading to logic gaps, syntax errors, and poor adherence to coding standards [85]. In contrast, open-source, code-specific models, trained on high-quality code and documentation, tend to produce more accurate and consistent outputs. Their open nature allows developers to inspect, secure, and customize them for specific needs, making these models more adaptable and reliable for real-world applications [86], [87]. Researchers have developed various code-focused LLMs, including CodeLlama [2], StarCoder [1], DeepSeek-Coder [3], CodeQwen [35], [33], and CodeStral [4]. These models generally follow one of two training paradigms: direct pretraining on code-specific corpora [88], [89], [43], or instruction tuning of foundation models [15], [90], [2]. Their training datasets include source

code repositories, text-code pairs, synthetic data, and mathematical or technical content [35], [3]. To evaluate the effectiveness of these models, a variety of standardized benchmarks and metrics have been developed, such as HumanEval, MBPP, and BigCodeBench [36], [38]. Diverse and large-scale training data enhances the performance of code LLMs but also increases the risk of exposing private information, making the DP essential for secure and regulation-compliant model training.

DP Dataset Synthesis. Previous works have used the DP framework to design secure dataset synthesis approaches across various data types, including tabular [19], [91], [92], image [18], [67], [93], [21], [94], text [64], [20], [65], [66], network trace [95], and others. Despite the different data types, most recent works have focused on a similar type of approach, sanitizing deep generative models using DP-SGD [22]. For example, Dockhorn et al. [67] trains diffusion models using DP-SGD for DP image synthesis. Yue et al. [20] propose training LLM using DP-SGD for DP text generation. Although code and text share some similar properties, unlike text, code has stricter syntax rules and structural dependencies. Therefore, directly applying existing DP text synthesis approaches to code will neglect such syntax dependencies.

VII. CONCLUSIONS

This paper presents the first DP code generation approach, providing a paradigm for learning specific knowledge from sensitive code datasets and generating code with high utility under the DP framework. PrivCode includes a two-stage training framework. The privacy-sanitizing stage introduces the PrivSA module for models on sensitive code datasets, incorporating adversarial code syntactic information to enhance the code generation capability of LLM, particularly for complex code generation. The utility-boosting stage uses high-quality, privacy-free datasets to train a premium model with no DP, avoiding additional performance degradation due to noise injection, and serves as the primary contributor to improving model utility. Comprehensive evaluations of PrivCode show that it outperforms prior DP fine-tuning methods across four benchmarks and achieves performance close to methods without any privacy protection mechanisms. We constructed the OSS-Instruct PII dataset and conducted

canary experiments to verify that PrivCode can effectively protect private information from the training dataset in real-world settings. Additionally, we performed ablation studies to demonstrate the importance of the PrivSA module and the two-stage paradigm in DP fine-tuning. Finally, we conduct hyper-parameter analysis experiments to evaluate the impact of different privacy budgets and max lambda. This work aims to promote the secure and responsible sharing of code datasets under DP, further advancing research in code LLMs.

ETHICS CONSIDERATIONS

We propose PrivCode, the first DP code synthesizer, providing a paradigm for learning specific knowledge from sensitive code datasets under the DP framework. All open-source datasets and models we used in this paper are publicly available and widely used in the community. All highlighted information we present in this paper is sourced from the open-source PII dataset [50] and synthetic canary samples. The PII dataset [50] explicitly states that the included PII data originates from open and permissively licensed GitHub repositories. The synthetic canary samples are generated using the open-source Llama-3.1-70B-Instruct [83], and the few-shot context also comes from the PII dataset [50]. To uphold the highest ethical standards and avoid any potential disclosure of personal private information, we further redact the presented public information using black blocks.

ACKNOWLEDGMENT

We thank anonymous reviewers for valuable comments. This paper was partially supported by NSF CNS-2220433 and a CCI research award. The opinions and findings expressed here are the authors' own and do not reflect the views of the National Science Foundation

REFERENCES

- [1] R. Li, L. B. Allal, Y. Zi *et al.*, “Starcoder: may the source be with you!” *arXiv preprint arXiv:2305.06161*, 2023.
- [2] B. Roziere, J. Gehring, F. Gloeckle *et al.*, “Code llama: Open foundation models for code,” *arXiv preprint arXiv:2308.12950*, 2023.
- [3] D. Guo, Q. Zhu, D. Yang *et al.*, “Deepseek-coder: When the large language model meets programming – the rise of code intelligence,” *arXiv preprint arXiv:2401.14196*, 2024.
- [4] MistralAI, “Codestral,” <https://mistral.ai/news/codestratal>, 2024.
- [5] A. Fan, B. Gokkaya, M. Harman, M. Lyubarskiy, S. Sengupta, S. Yoo, and J. M. Zhang, “Large language models for software engineering: Survey and open problems,” in *2023 IEEE/ACM International Conference on Software Engineering: Future of Software Engineering (ICSE-FoSE)*. IEEE, 2023, pp. 31–53.
- [6] B. Min, H. Ross, E. Sulem, A. P. B. Veyseh, T. H. Nguyen, O. Sainz, E. Aguirre, I. Heintz, and D. Roth, “Recent advances in natural language processing via large pre-trained language models: A survey,” *ACM Computing Surveys*, vol. 56, no. 2, pp. 1–40, 2023.
- [7] R. A. Husein, H. Aburajouh, and C. Catal, “Large language models for code completion: A systematic literature review,” *Computer Standards & Interfaces*, vol. 92, p. 103917, 2025.
- [8] M. A. Tashtoush, Y. Wardat, F. Aloufi, and O. Taani, “The effect of a training program based on tims to developing the levels of habits of mind and mathematical reasoning skills among pre-service mathematics teachers,” *EURASIA Journal of Mathematics, Science and Technology Education*, vol. 18, no. 11, p. em2182, 2022.
- [9] Y. Yang, X. Hu, Z. Gao, J. Chen, C. Ni, X. Xia, and D. Lo, “Federated learning for software engineering: A case study of code clone detection and defect prediction,” *IEEE Transactions on Software Engineering*, 2024.
- [10] N. Carlini, D. Ippolito, M. Jagielski, K. Lee, F. Tramer, and C. Zhang, “Quantifying memorization across neural language models,” in *The Eleventh International Conference on Learning Representations*, 2022.
- [11] N. Carlini, F. Tramer, E. Wallace *et al.*, “Extracting training data from large language models,” in *30th USENIX security symposium (USENIX Security 21)*, 2021, pp. 2633–2650.
- [12] M. Nasr, N. Carlini, J. Hayase *et al.*, “Scalable extraction of training data from (production) language models,” *arXiv preprint arXiv:2311.17035*, 2023.
- [13] A. Radford, J. Wu, R. Child *et al.*, “Language models are unsupervised multitask learners,” *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [14] L. Niu, S. Mirza, Z. Maradni, and C. Pöpper, “{CodexLeaks}: Privacy leaks from code generation language models in {GitHub} copilot,” in *32nd USENIX Security Symposium*, 2023, pp. 2133–2150.
- [15] M. Chen, J. Tworek, H. Jun *et al.*, “Evaluating large language models trained on code,” *arXiv preprint arXiv:2107.03374*, 2021.
- [16] H. Pearce, B. Ahmad, B. Tan, B. Dolan-Gavitt, and R. Karri, “Asleep at the keyboard? assessing the security of github copilot’s code contributions,” *2022 IEEE Symposium on Security and Privacy (SP)*, pp. 754–768, 2022.
- [17] Github, “Github copilot - your ai pair programmer,” <https://copilot.github.com/>, 2022.
- [18] K. Li, C. Gong, Z. Li, Y. Zhao, X. Hou, and T. Wang, “{PrivImage}: Differentially private synthetic image generation using diffusion models with {Semantic-Aware} pretraining,” in *33rd USENIX Security Symposium (USENIX Security 24)*, 2024, pp. 4837–4854.
- [19] K. Cai, X. Lei, J. Wei, and X. Xiao, “Data synthesis via differentially private markov random fields,” *Proc. VLDB Endow.*, vol. 14, no. 11, p. 2190–2202, 2021.
- [20] X. Yue, H. A. Inan, X. Li, G. Kumar, J. McAnallen, H. Shahari, H. Sun, D. Levitan, and R. Sim, “Synthetic text generation with differential privacy: A simple and practical recipe,” *arXiv preprint arXiv:2210.14348*, 2022.
- [21] C. Gong, K. Li, Z. Lin, and T. Wang, “Dpimagebench: A unified benchmark for differentially private image synthesis,” *arXiv preprint arXiv:2503.14681*, 2025.
- [22] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, “Deep learning with differential privacy,” in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 308–318.
- [23] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, “Lora: Low-rank adaptation of large language models,” *arXiv preprint arXiv:2106.09685*, 2021.
- [24] M. Singh, J. Cambronero, S. Gulwani, V. Le, C. Negreanu, and G. Verbruggen, “Codefusion: A pre-trained diffusion model for code generation,” in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2023, pp. 11 697–11 708.
- [25] Y. Zhu, Y. Zhang, H. Yang, and F. Wang, “Gancoder: an automatic natural language-to-programming language translation approach based on gan,” in *CCF International Conference on Natural Language Processing and Chinese Computing*. Springer, 2019, pp. 529–539.
- [26] T. Dockhorn, T. Cao, A. Vahdat, and K. Kreis, “Differentially private diffusion models,” *arXiv preprint arXiv:2210.09929*, 2022.
- [27] W. Ma, S. Liu, M. Zhao, X. Xie, W. Wang, Q. Hu, J. Zhang, and Y. Liu, “Unveiling code pre-trained models: Investigating syntax and semantics capacities,” *ACM Transactions on Software Engineering and Methodology*, vol. 33, no. 7, pp. 1–29, 2024.
- [28] C. Hou, A. Shrivastava, H. Zhan, R. Conway, T. Le, A. Sagar, G. Fanti, and D. Lazar, “Pre-text: Training language models on private federated data in the age of llms,” *arXiv preprint arXiv:2406.02958*, 2024.
- [29] C. Dwork, A. Roth *et al.*, “The algorithmic foundations of differential privacy,” *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.
- [30] Y. Li, D. Choi, J. Chung *et al.*, “Competition-level code generation with alphacode,” *Science*, vol. 378, no. 6624, p. 1092–1097, Dec. 2022. [Online]. Available: <http://dx.doi.org/10.1126/science.abq1158>
- [31] M. Allamanis, S. Panthaplectak, and P. Yin, “Unsupervised evaluation of code llms with round-trip correctness,” *arXiv preprint arXiv:2402.08699*, 2024.
- [32] L. Ben Allal, N. Muennighoff, L. Kumar Umapathi *et al.*, “A framework for the evaluation of code generation models,” <https://github.com/bigcode-project/bigcode-evaluation-harness>, 2022.
- [33] B. Hui, J. Yang, Z. Cui *et al.*, “Qwen2. 5-coder technical report,” *arXiv preprint arXiv:2409.12186*, 2024.

[34] C. Team, H. Zhao, J. Hui *et al.*, “Codegemma: Open code models based on gemma,” *arXiv preprint arXiv:2406.11409*, 2024.

[35] Q. Team, “Code with codeqwen1.5,” April 2024. [Online]. Available: <https://qwenlm.github.io/blog/codeqwen1.5/>

[36] J. Austin, A. Odena, M. Nye *et al.*, “Program synthesis with large language models,” *arXiv preprint arXiv:2108.07732*, 2021.

[37] J. Liu, C. S. Xia, Y. Wang, and L. Zhang, “Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.

[38] T. Y. Zhuo, M. C. Vu, J. Chim *et al.*, “Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions,” *arXiv preprint arXiv:2406.15877*, 2024.

[39] J. Wei, M. Bosma, V. Y. Zhao *et al.*, “Finetuned language models are zero-shot learners,” *arXiv preprint arXiv:2109.01652*, 2021.

[40] M. Bruch, M. Monperrus, and M. Mezini, “Learning from examples to improve code completion systems,” in *Proceedings of the 7th joint meeting of the European software engineering conference and the ACM SIGSOFT symposium on the foundations of software engineering*, 2009, pp. 213–222.

[41] Anthropic, “Claude 4,” <https://www.anthropic.com/news/clause-4>, 2025.

[42] OpenAI, “Gpt-4o,” <https://openai.com/index/hello-gpt-4o>, 2024.

[43] Y. Wang, W. Wang, S. Joty, and S. C. H. Hoi, “Codet5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation,” *arXiv preprint arXiv:2109.00859*, 2021.

[44] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to sequence learning with neural networks,” *arXiv preprint arXiv:1409.3215*, 2014.

[45] P. Koehn, “Statistical significance tests for machine translation evaluation,” in *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, Jul. 2004, pp. 388–395.

[46] A. Fan, M. Lewis, and Y. Dauphin, “Hierarchical neural story generation,” *arXiv preprint arXiv:1805.04833*, 2018.

[47] S. Aho, R. Sethi, and D. Jeffrey, “Ullman, compilers: Principles, techniques, and tools,” 1986.

[48] I. D. Baxter, A. Yahin, L. Moura *et al.*, “Clone detection using abstract syntax trees,” in *Proceedings. International Conference on Software Maintenance*, 1998, pp. 368–377.

[49] D. Ippolito, F. Tramèr, M. Nasr *et al.*, “Preventing verbatim memorization in language models gives a false sense of privacy,” *arXiv preprint arXiv:2210.17546*, 2022.

[50] L. B. Allal, R. Li, D. Kocetkov *et al.*, “Santacoder: don’t reach for the stars!” *arXiv preprint arXiv:2301.03988*, 2023.

[51] D. Kocetkov, R. Li, L. B. Allal *et al.*, “The stack: 3 tb of permissively licensed source code,” *arXiv preprint arXiv:2211.15533*, 2022.

[52] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Artificial intelligence and statistics*, 2017, pp. 1273–1282.

[53] J. Lebensold, M. Sanjabi, P. Astolfi *et al.*, “Dp-rdm: Adapting diffusion models to private domains without fine-tuning,” *arXiv preprint arXiv:2403.14421*, 2024.

[54] S. Romansky, C. Chen, B. Malhotra, and A. Hindle, “Sorcerer’s apprentice and the study of code snippet migration,” *arXiv preprint arXiv:1808.00106*, 2018.

[55] C. Dwork, F. McSherry, K. Nissim, and A. Smith, “Calibrating noise to sensitivity in private data analysis,” in *Theory of Cryptography: Third Theory of Cryptography Conference*, 2006, pp. 265–284.

[56] F. Mireshghallah, F. Tramèr, H. Brown *et al.*, “What does it mean for a language model to preserve privacy?” in *FACCT*, 2022.

[57] A. Sinha, T. Mesnard, R. McKenna *et al.*, “Vaultgemma: A differentially private gemma model,” 2025. [Online]. Available: <https://arxiv.org/abs/2510.15001>

[58] I. Mironov, K. Talwar, and L. Zhang, “R\’enyi differential privacy of the sampled gaussian mechanism,” *arXiv preprint arXiv:1908.10530*, 2019.

[59] X. Zhang, J. Zhao, and Y. LeCun, “Character-level convolutional networks for text classification,” *Advances in neural information processing systems*, 2015.

[60] Y. Wei, Z. Wang, J. Liu, Y. Ding, and L. Zhang, “Magicoder: Empowering code generation with oss-instruct,” *Forty-first International Conference on Machine Learning*, 2024.

[61] C. E. Shannon, “A mathematical theory of communication,” *The Bell system technical journal*, vol. 27, no. 3, pp. 379–423, 1948.

[62] C. Xie, Z. Lin, A. Backurs *et al.*, “Differentially private synthetic data via foundation model apis 2: Text,” *Forty-first International Conference on Machine Learning*, 2024.

[63] S. Utpala, S. Hooker, and P. Y. Chen, “Locally differentially private document generation using zero shot prompting,” *arXiv preprint arXiv:2310.16111*, 2023.

[64] J. Hong, J. T. Wang, C. Zhang, Z. Li, B. Li, and Z. Wang, “Dp-opt: Make large language model your privacy-preserving prompt engineer,” *arXiv preprint arXiv:2312.03724*, 2023.

[65] D. Yu, P. Kairouz, S. Oh, and Z. Xu, “Privacy-preserving instructions for aligning large language models,” *preprint arXiv:2402.13659*, 2024.

[66] A. G. Carranza, R. Farahani, N. Ponomareva *et al.*, “Synthetic query generation for privacy-preserving deep retrieval systems using differentially private language models,” *arXiv preprint arXiv:2305.05973*, 2023.

[67] T. Dockhorn, T. Cao, A. Vahdat *et al.*, “Differentially private diffusion models,” *Transactions on Machine Learning Research*, 2023.

[68] P. J. Guo and D. Engler, “{CDE}: Using system call interposition to automatically create portable software packages,” in *2011 USENIX Annual Technical Conference (USENIX ATC 11)*, 2011.

[69] B. Yee, D. Sehr, G. Dardyk *et al.*, “Native client: A sandbox for portable, untrusted x86 native code,” *Communications of the ACM*, vol. 53, no. 1, pp. 91–99, 2010.

[70] N. Muennighoff, Q. Liu, A. R. Zebaze *et al.*, “Octopack: Instruction tuning code large language models,” in *The Twelfth International Conference on Learning Representations*, 2024.

[71] T. Zhang, V. Kishore, F. Wu *et al.*, “Bertscore: Evaluating text generation with bert,” *arXiv preprint arXiv:1904.09675*, 2019.

[72] D. Yu, S. Naik, A. Backurs, S. Gopi, H. A. Inan, G. Kamath, J. Kulkarni, Y. T. Lee, A. Manoel, L. Wutschitz *et al.*, “Differentially private fine-tuning of language models,” *arXiv preprint arXiv:2110.06500*, 2021.

[73] J. Xu, K. Saravanan, R. van Dalen, H. Mehmod, D. Tuckey, and M. Ozay, “Dp-dylora: Fine-tuning transformer-based models on-device under differentially private federated learning using dynamic low-rank adaptation,” *arXiv preprint arXiv:2405.06368*, 2024.

[74] N. Houlsby, A. Giurgiu, S. Jastrzebski *et al.*, “Parameter-efficient transfer learning for nlp,” in *International conference on machine learning*, 2019, pp. 2790–2799.

[75] W. Shi, R. Shea, S. Chen, C. Zhang, R. Jia, and Z. Yu, “Just fine-tune twice: Selective differential privacy for large language models,” *arXiv preprint arXiv:2204.07667*, 2022.

[76] Q. Zheng, X. Xia, X. Zou *et al.*, “Codegeex: A pre-trained model for code generation with multilingual benchmarking on humaneval-x,” in *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2023, pp. 5673–5684.

[77] Z. Liu, C. Gong, T. Y. Zhuo, K. Li, W. Yu, M. Fredrikson, and T. Wang, “Privcode: When code generation meets differential privacy,” *arXiv preprint arXiv:2512.05459*, 2025.

[78] N. Kandpal, E. Wallace, and C. Raffel, “Deduplicating training data mitigates privacy risks in language models,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 10 697–10 707.

[79] D. P. W. Party, “Opinion 05/2014 on anonymisation techniques,” *Diunduh Dari*, 2014.

[80] A. Grattafiori, A. Dubey, A. Jauhri *et al.*, “The llama 3 herd of models,” *arXiv preprint arXiv:2407.21783*, 2024.

[81] J. Rasley, S. Rajbhandari, O. Ruwase, and Y. He, “Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters,” in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, p. 3505–3506.

[82] Y. Yang, S. Mishra, J. Chiang, and B. Mirzaoleiman, “Smalltolarge (s2l): Scalable data selection for fine-tuning large language models by summarizing training trajectories of small models,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 83 465–83 496, 2024.

[83] D. Patterson, J. Gonzalez, U. Hölzle *et al.*, “The carbon footprint of machine learning training will plateau, then shrink,” *Computer*, vol. 55, no. 7, pp. 18–28, 2022.

[84] K. Lee, D. Ippolito, A. Nystrom *et al.*, “Deduplicating training data makes language models better,” *arXiv preprint arXiv:2107.06499*, 2021.

[85] F. F. Xu, U. Alon, G. Neubig *et al.*, “A systematic evaluation of large language models of code,” in *Proceedings of the 6th ACM SIGPLAN International Symposium on Machine Programming*, 2022, pp. 1–10.

[86] M. Du, A. T. Luu, B. Ji, Q. Liu, and S.-K. Ng, “Mercury: A code efficiency benchmark for code large language models,” in *The Thirty-eighth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024.

[87] J. He, M. Vero, G. Krasnopska, and M. Vechev, “Instruction tuning for secure code generation,” *arXiv preprint arXiv:2402.09497*, 2024.

[88] E. Nijkamp, B. Pang, H. Hayashi *et al.*, “Codegen: An open large language model for code with multi-turn program synthesis,” *arXiv preprint arXiv:2203.13474*, 2022.

[89] D. Fried, A. Aghajanyan, J. Lin *et al.*, “Incoder: A generative model for code infilling and synthesis,” *arXiv preprint arXiv:2204.05999*, 2022.

[90] A. Lozhkov, R. Li, L. B. Allal *et al.*, “Starcoder 2 and the stack v2: The next generation,” *arXiv preprint arXiv:2402.19173*, 2024.

[91] Z. Zhang, T. Wang, N. Li *et al.*, “{PrivSyn}: Differentially private data synthesis,” in *USENIX Security Symposium*, 2021, pp. 929–946.

[92] R. McKenna, B. Mullins, D. Sheldon, and G. Miklau, “AIM: an adaptive and iterative mechanism for differentially private synthetic data,” *Proc. VLDB Endow.*, vol. 15, no. 11, pp. 2599–2612, 2022.

[93] Z. Lin, S. Gopi, J. Kulkarni, H. Nori, and S. Yekhanin, “Differentially private synthetic data via foundation model APIs 1: Images,” in *The Twelfth International Conference on Learning Representations*, 2024.

[94] K. Li, C. Gong, X. Li *et al.*, “From easy to hard: Building a shortcut for differentially private image synthesis,” in *2025 IEEE Symposium on Security and Privacy (SP)*, 2025, pp. 3656–3674.

[95] D. Sun, J. Q. Chen, C. Gong, T. Wang, and Z. Li, “Netdpsyn: synthesizing network traces under differential privacy,” in *Proceedings of the 2024 ACM on Internet Measurement Conference*, 2024, pp. 545–554.

[96] W. Kwon, Z. Li, S. Zhuang *et al.*, “Efficient memory management for large language model serving with pagedattention,” in *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.

[97] Z. Wang, Z. Zhou, D. Song, Y. Huang, S. Chen, L. Ma, and T. Zhang, “Where do large language models fail when generating code?” *arXiv e-prints*, pp. arXiv–2406, 2024.

[98] W. Wang, Z. Li, D. Lian, C. Ma, L. Song, and Y. Wei, “Mitigating the language mismatch and repetition issues in llm-based machine translation via model editing,” *arXiv preprint arXiv:2410.07054*, 2024.

APPENDIX

Due to space limitations, please refer to the full version [77] for additional appendices.

A. Details of Studied Datasets

Magicoder-OSS-Instruct-75K. This dataset is generated by Magicoder [60] using its OSS-Instruct method for instruction fine-tuning. It contains a large number of high-quality task-code instruction pairs. OSS-Instruct is a prompt engineering method for open-source code that utilizes a vast amount of code from open-source software (OSS) repositories. Constructing a carefully designed prompt automatically generates useful instructions or task descriptions. The primary goal is to extract high-quality data from real-world codebases for code generation tasks while retaining the core information.

OSS-Instruct PII Dataset. The PII dataset [50] is an annotated dataset for PIIs in code. The target entities include: Names, Usernames, Emails, IP addresses, Keys, Passwords, and IDs. The annotation process involved 1,399 crowdworkers from 35 countries using Toloka. The dataset consists of 12,099 samples, each approximately 50 lines of code, in 31 programming languages. The PII dataset was constructed by manually annotating the following entities on a small portion of The Stack [51] by 12 members of the BigCode community: Names, Emails, Usernames, Passwords, IP addresses, API Keys, and SSH Keys. We provide more details of the dataset in Appendix A of our full version [77].

B. Details of Baselines

This section introduces the studied baselines, including NonDPFT, JFT, and DPFT. We provide more details of baselines in Appendix B [77].

TABLE VIII: The GPU memory consumption and running time of each experiment stage of PrivCode and baselines. “-” means the current stage only cost CPU memory.

| Method | Stage | GPU Memory | Running Time |
|----------|--------------------------------|------------|--------------|
| PrivCode | Privacy-sanitizing Fine-tuning | 24.4GB | 2.37h |
| | Privacy-free Data Synthesis | 15.8GB | 0.39h |
| | Execution Validation | – | 0.25h |
| | Round-trip Validation | 78.7GB | 0.45h |
| | Utility-boosting Fine-tuning | 32.4GB | 1.13h |
| JFT | The First Non-DP Fine-tuning | 29.7GB | 2.11h |
| | The Second DP Fine-tuning | 37.3GB | 1.74h |
| DPFT | DP Fine-tuning Stage | 38.2GB | 2.08h |

C. DP-SGD Hyper-parameter Settings

This section provides a detailed description of the implementation of the DP-SGD training process in each method under target $\epsilon = 4.0$. We follow the FastDP repository,² to conduct our experiments. The key hyper-parameters are detailed in Table X. In PrivCode, we set a smaller maximum training step because the junior model (Qwen2.5-Coder-1.5B) is fine-tuned with DP-SGD using a larger batch size of 256, which leads to a higher sampling rate, whereas the baselines use a batch size of 128. The dataset size for DPFT is kept consistent with that used in the utility-boosting stage of PrivCode to ensure fairness. The noise scale σ is computed using the standard privacy analysis function provided by FastDP. For each DP-SGD training experiment, we set the clipping norm C to 1.0, δ to 1e-5, and the accountant type to RDP, following the default configuration of FastDP.

D. Consumption

This section provides a detailed cost analysis that compares PrivCode with baselines. Table VIII shows the GPU memory consumption and running time of each process of PrivCode and baselines. We conduct our experiments on four NVIDIA GeForce A6000 GPUs with 48 GB of memory, except that the round-trip validation of PrivCode is performed on a server equipped with four NVIDIA A100 GPUs with 80 GB of memory. We provide more discussions in Appendix E [77].

Regarding running time consumption, we acknowledge that PrivCode requires additional running time due to its multi-stage workflow. However, we utilize DeepSpeed [81] to accelerate the fine-tuning stages, vLLM [96] to parallelize inference in the privacy-free data synthesis and round-trip validation stages, and CPU parallel processing in the execution validation, thereby reducing the total running time of PrivCode to an acceptable range compared to the baselines. Additionally, the privacy-sanitizing fine-tuning stage incurs extra runtime overhead due to the integration of the PrivSA module. We believe that PrivCode is a practical and effective approach that achieves a favorable consumption–utility trade-off.

E. Details of Execution Validation

We provide a more detailed explanation and implementation of execution validation, and report the statistical breakdowns

²<https://github.com/awslabs/fast-differential-privacy>

TABLE IX: The inference prompt is similar in form or semantics to the training prompt shown in Table I. Code snippets are generated by **PrivCode** under $\epsilon = 1, 4, 10$ and by **NonDPFT** under $\epsilon = \infty$. The private information is highlighted.

| Method | Text |
|----------|---|
| Prompt | Write a Python function that processes a list of dictionaries. Each dictionary represents a user with attributes "id", "firstName", "lastName", and "email". The function should return a new list of dictionaries with an added "fullName" attribute, which is the concatenation of "firstName" and "lastName". Ensure the output matches the following format: <code>[{"id": 1, "firstName": "FirstName", "lastName": "LastName", "email": "example@example.com", "fullName": "FirstName LastName"}, ...]</code> |
| NonDPFT | <code>\'\'\'python\ndef process_users(users):\n for user in users:\n user["fullName"]\n = f"{user['firstName']} {user['lastName']}"\n return users\n\nusers = [\n {"id": 1, "firstName": "Eli", "lastName": "Gerry", "email": "eli@tryasity.com", "id": 2, "firstName": "Raled", "lastName": "Palk", "email": "ra@rk@asity.com", "id": 3, "firstName": "Ollin", "lastName": "Ollin", "email": "owen@asity.com", "id": 4, "firstName": "Brina", "lastName": "Galilia", "email": "brina@asity.com"}\n\nupdated_users = process_users(users)\nprint(updated_users)\'\'\'</code> |
| | <code>\'\'\'python\ndef filter_user_attributes(users):\n filtered_users = []\n for user in users:\n filtered_user = {"id": user['id'], "firstName": user['firstName'], "lastName": user['lastName'], "email": user['email'], "fullName": user['firstName'] + ' ' + user['lastName']}\n filtered_users.append(filtered_user)\n return filtered_users\n\'\'\'</code> |
| | <code>\'\'\'python\ndef process_users(users):\n return [\n {"id": user["id"], "firstName": user["firstName"], "lastName": user["lastName"], "email": user["email"], "fullName": f"{user['firstName']} {user['lastName']}"}\n for user in users\n]\n\'\'\'</code> |
| PrivCode | <code>\'\'\'python\ndef add_full_name(users):\n for user in users:\n user["fullName"] = f"{user['firstName']} {user['lastName']}"\n return users\n\'\'\'</code> |

TABLE X: DP-SGD hyper-parameter settings under target $\epsilon = 4.0$. The sampling rate q is computed by the dataset size and batch size. We use AdamW optimizer with a learning rate of 5e-6.

| Method | Dataset Size | Sampling Rate q | Max Step | Clipping Norm C | Noise Scale σ | δ | Accountant | Resulting ϵ |
|----------|--------------|-------------------|----------|-------------------|----------------------|----------|------------|----------------------|
| DPFT | 15310 | 0.0083 | 2000 | 1.0 | 0.77 | 1e-5 | RDP | 3.98 |
| JFT | 19551 | 0.0262 | 1000 | 1.0 | 0.69 | 1e-5 | RDP | 3.97 |
| PrivCode | 19551 | 0.0131 | 100 | 1.0 | 0.63 | 1e-5 | RDP | 3.97 |

TABLE XI: The number / distribution of instances for each failure cause classification in the execution validation under the settings of Section V-A. The total number of instances to be execution-filtered for each model is the same.

| Model | Environment Error | Compile Error | Runtime Error | Language Mismatch | Others |
|------------------|-------------------|---------------|----------------|-------------------|-------------|
| DS-Coder-6.7B | 10232 / 25.77% | 4678 / 11.79% | 7206 / 18.16% | 17189 / 43.33% | 387 / 0.98% |
| Qwen2.5-Coder-7B | 12162 / 30.64% | 3261 / 8.22% | 8756 / 22.06% | 15330 / 38.62% | 183 / 0.46% |
| CodeGemma-7B | 13580 / 34.20% | 2494 / 6.27% | 10106 / 25.47% | 13079 / 32.91% | 433 / 1.09% |
| CodeQwen1.5-7B | 11243 / 28.32% | 5148 / 12.97% | 9581 / 24.12% | 13386 / 33.70% | 334 / 0.84% |

of practical run-time failure causes, along with analysis in the full version of our paper [77].

We report the statistical breakdowns of practical run-time failure causes in the execution validation under the experiment settings of Section V-A. Table XI shows that language mismatch is the dominant cause of execution failures, indicating that the model produces code containing mixed programming-language syntax or code that is entirely inconsistent with the specified language [97], [98]. Execution validation is essential for filtering out the synthetic low-quality code. We provide more details in Appendix F of our full version paper [77].

F. Robustness

We conduct a robustness study of **PrivCode**'s training. We set random seeds in the fine-tuning steps, privacy-free code snippets synthesis step, and round-trip validation step of **PrivCode**, along with training steps of baselines, to

demonstrate the stable utility superiority under randomness. Table XII reports the pass@1 scores (mean \pm standard deviation) of **PrivCode** and the baselines under $\epsilon = 4$ when using Qwen2.5-Coder-7B as the base model. Even under different random seeds, **PrivCode** consistently outperforms the other privacy-preserving baselines across all benchmarks, achieving improvements of up to 11.4% on instruction-following benchmarks and 10.3% on code-completion benchmarks.

G. Abalition Study

Figure 9 of our full version of paper [77] presents the examples of code generated in the HumanEval benchmark using the same prompt, where models are triggered to generate more complex code. Figure 7 [77] presents examples of relatively simple code generated in the MBPP benchmark. We use DS-Coder-6.7B [3] as the evaluated LLM. While both **PrivCode** and NonASTPrivCode pass the MBPP test

TABLE XII: Pass@1 score (mean \pm standard deviation) of PrivCode and baselines under $\epsilon = 4$ trained using Qwen2.5-Coder-7B as premium model. The bolded data represents the best score, and the gray shaded area indicates PrivCode.

| Method | HumanEval | | MBPP | | HumanEval | | MBPP | |
|---------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|
| | HE | HE+ | MBPP | MBPP+ | HE | HE+ | MBPP | MBPP+ |
| | Instruct | | | | Complete | | | |
| NonDPFT | 77.6 \pm 3.2 | 74.1 \pm 2.5 | 75.1 \pm 8.3 | 58.4 \pm 4.9 | 69.7 \pm 3.6 | 57.2 \pm 5.3 | 76.9 \pm 5.5 | 59.8 \pm 9.4 |
| DPFT | 62.8 \pm 4.3 | 54.0 \pm 7.4 | 57.8 \pm 2.2 | 46.9 \pm 6.5 | 34.7 \pm 3.8 | 28.0 \pm 0.4 | 59.5 \pm 3.9 | 46.3 \pm 7.0 |
| JFT | 64.3 \pm 6.1 | 54.9 \pm 3.7 | 66.7 \pm 5.5 | 61.5 \pm 4.1 | 35.9 \pm 5.6 | 36.1 \pm 6.2 | 69.1 \pm 2.8 | 54.4 \pm 4.8 |
| Ours | 67.3 \pm 4.3 | 60.6 \pm 3.8 | 78.1 \pm 5.6 | 66.5 \pm 7.5 | 42.8 \pm 1.7 | 40.3 \pm 4.2 | 75.4 \pm 4.1 | 64.7 \pm 6.2 |

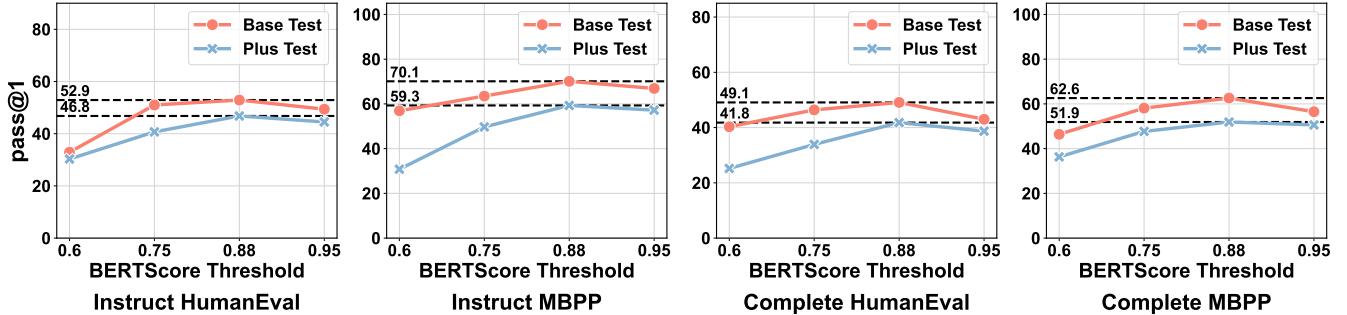


Fig. 7: Average pass@1 of four LLMs trained using PrivCode, evaluated on HumanEval, MBPP, and EvalPlus benchmarks. “Base Test” is HumanEval or MBPP, while “Plus Test” means HumanEval+ and MBPP+. The pass@1 score varies with different BERTScore threshold τ_s . The dashed lines indicate the top value.

case, only PrivCode successfully passes the HumanEval test case. Examples show that PrivCode has a stronger ability to generate more complex code, with higher structural and syntactic accuracy compared to the NonASTPrivCode.

H. Analysis of Validation Filters

We provide a detailed analysis of the acceptance rates of the execution and round-trip validation filters, the sensitivity analysis of the BERTScore threshold in the round-trip filter, and a taxonomy contrasting syntax- and semantic-level failures.

Acceptance Rates of Validation Filters. Under the experimental setup of Section V-A, we generate 55,500 synthetic samples. After applying the execution validation, 31,644 samples remain, with 57% acceptance rate, indicating that a substantial portion of model outputs contain syntax errors, runtime failures, missing dependencies, or other execution-level issues, referring to Appendix E. Subsequently, the round-trip validation further filters these 31,644 samples and retains 15,311, under the BERTScore threshold of 0.88, with 48% acceptance rate, resulting in only 27.6% of the total synthetic samples surviving both filters. This demonstrates that both stages of validation are necessary to ensure syntactic correctness and semantic fidelity.

Sensitivity to the BERTScore Threshold τ_s . We conduct hyper-parameter study against the round-trip validator under thresholds $\tau_s \in \{0.60, 0.75, 0.88, 0.95\}$ to understand its impact on semantic filtering strictness and final utility of PrivCode. The acceptance rate of round-trip validation directly reflects the sensitivity to semantic filtering strictness. We observe that lower thresholds substantially increase acceptance but introduce semantic drift: $\tau_s = 0.60$ admits 25,491 samples (46% acceptance), and $\tau_s = 0.75$ retains 19,842 samples (36% acceptance), yet both settings allow many semantically

inconsistent or loosely aligned outputs. In contrast, a very high threshold such as $\tau_s = 0.95$ is overly strict and keeps only 5,985 samples (11% acceptance), sharply reducing usable data. The mid-range threshold $\tau_s = 0.88$, which retains 15,311 samples (28% acceptance), achieves the best balance between preserving dataset scale and maintaining semantic fidelity.

Figure 7 shows that the final utility of PrivCode initially increases with τ_s and then decreases beyond an optimal point. Under the studied thresholds $\{0.60, 0.75, 0.88, 0.95\}$, $\tau_s = 0.88$ achieves the best trade-off. While a high threshold such as $\tau_s = 0.95$ ensures semantic correctness of the filtered data, it reduces diversity and dataset size, yielding high pass@1 scores across four benchmark tests $\{49.4, 66.9, 43.0, 56.5\}$. Conversely, a low threshold such as $\tau_s = 0.60$ retains amount low-quality samples with syntax errors, resulting in substantial performance drops $\{32.9, 56.9, 40.3, 46.4\}$. We clarify that excessively high thresholds lead to maintain samples lacking diversity and then increase the risk of overfitting, reducing generalization performance.

Filter Failure Taxonomy. The execution validation primarily removes *syntax-level* failures, including compile errors, runtime exceptions, missing dependencies, and language mismatches (as detailed in Appendix E). In contrast, round-trip validation targets *semantic-level* deviations rather than syntactic failures. It captures cases where the generated code diverges from the prompt in functionality, reproduces only part of the intended behavior, or is loosely related to the instruction, resulting in low round-trip similarity. Unlike execution validation, these failures reflect shortcomings in semantic fidelity rather than syntactic correctness.

We present a case study showing how PrivCode can avoid leaking private information from the training dataset in Table IX, and provide more details in Appendix K [77].

Artifact Appendix

Abstract. This is the official implementation of paper PrivCode: When code generation Meets Differential Privacy. This repository contains PyTorch training code and evaluation code. PrivCode is a Differential Privacy (DP) code generation tool that leverages DP to generate synthetic code, enabling organizations to share and use code LLMs without privacy concerns. Our proposed artifacts support this paper in the following ways: (1) *Training Framework*: Implements DP-based optimization strategies for code generation, ensuring that generated code does not leak sensitive information from the original dataset. (2) *Evaluation Pipeline*: Provides the leakage rate metric of the canary experiment for privacy and the pass@1 score metric for utility, enabling reproducible experiments. (3) *Synthetic Code Filtering Module*: Screens generated code for quality and compliance, ensuring that only useful and privacy-safe samples are retained. We have released the source code and included detailed documentation.

A. Description & Requirements

1) *How to access*: Our implementation can be accessed from the public GitHub repository. Please refer to the link: <https://github.com/Liuzzyg/PrivCode>.

2) *Hardware dependencies*: Training of all LLMs is conducted on a server configured with Python 3.11.0, equipped with four NVIDIA GeForce A6000 GPUs and 48GB of memory. The round-trip validation is conducted on a server equipped with four NVIDIA A100 GPUs. Actually, our experiments only require at least one A100 GPU with more than 80GB of memory.

3) *Software dependencies*: Our artifact requires the installation of DeepSpeed³, Peft⁴, Bitsandbyte⁵, and vLLM⁶. Otherwise, we encourage performing the execution validation in our Docker image⁷, which can be pulled from Docker Hub.

B. Artifact Installation & Configuration

For installation, please refer to the link: <https://github.com/Liuzzyg/PrivCode/README.md>. We elaborate on each step in this "README.md" file. Training of all LLMs is conducted on a server configured with Python 3.11.0. We provide a step-by-step guide on how to install our repository using Anaconda. We conduct experiments on Magicoder-OSS-Instruct-75K⁸ and OSS-Instruct PII dataset⁹, using Qwen2.5-Coder-1.5B¹⁰, deepseek-coder-6.7b-base¹¹, Qwen2.5-Coder-7B¹², codegemma-7b¹³, and Qwen/CodeQwen1.5-7B¹⁴. This

³<https://github.com/deepspeedai/DeepSpeed>

⁴<https://github.com/huggingface/peft>

⁵<https://github.com/bitsandbytes-foundation/bitsandbytes>

⁶<https://github.com/vllm-project/vllm>

⁷<https://hub.docker.com/repository/docker/liuzhengy/privcode-execution-filter>

⁸<https://huggingface.co/datasets/ise-uiuc/Magicoder-OSS-Instruct-75K>

⁹<https://huggingface.co/datasets/bigcode/bigcode-pii-dataset>

¹⁰<https://huggingface.co/Qwen/Qwen2.5-Coder-1.5B>

¹¹<https://huggingface.co/deepseek-ai/deepseek-coder-6.7b-base>

¹²<https://huggingface.co/Qwen/Qwen2.5-Coder-7B>

¹³<https://huggingface.co/google/codegemma-7b>

¹⁴<https://huggingface.co/Qwen/CodeQwen1.5-7B>

repository can be easily extended to other datasets and models downloaded from huggingface. This "README.md" file provides scripts for replicating our experiments.

C. Experiment Workflow

We describe the experimental workflow of PrivCode corresponding to the results in Table 3, Table 4, and Table 5 of our paper. All scripts are provided in this repository, and the following instructions reproduce the experiments step by step.

• Utility Experiment (Table 3).

– *Step 1: Privacy-sanitizing Stage.* To perform fine-tuning with the Privacy-free Syntax-Aware (PrivSA) module, please run "bash scripts-run-finetune/privcode_privacy_sanitizing.sh" and "python scripts-run-merge-peft/privcode_privacy_sanitizing.py". The privacy budget can be controlled by adjusting the variable "TARGET_EPSILONS".

– *Step 2: Utility-boosting Stage.* To generate privacy-free data, please run "bash data/private_syn/run_generate.sh". An example saving path of the generated data is "data/private_syn/Qwen2.5-Coder-1.5B/dp4_lambda1000to0.1_alpha0.01.jsonl".

To apply the execution filter, please pull and run the Docker image "liuzhengy/privcode-execution-filter:latest", copy the generated data into the container, and run "bash data/private_syn/run_clean_data.sh" inside it. After execution filtering, copy the cleaned file back to the local directory.

For round-trip filtering, please run "bash data/private_syn/run_rt_test_prompt.sh".

For fine-tuning without DP, please run "bash scripts-run-finetune/privcode_utility_boosting.sh" and "python scripts-run-merge-peft/privcode_utility_boosting.py".

– *Step 3: Utility Evaluation.* To evaluate the performance on the EvalPlus benchmark, please run "bash eval-utility/eval-evalplus/run_evalplus_privcode.sh". To evaluate on the BigCodeBench benchmark, please run "bash eval-utility/eval-bcb/run_bigcodebench_privcode.sh".

• Canary Experiment (Table 4).

– The training procedure follows the same workflow as the Utility Experiment described above, except that canary-injected OSS-Instruct PII datasets are used as sensitive datasets. You can obtain these datasets from <https://huggingface.co/datasets/ZhengLiu33/OSS-Instruct-PII-dataset> or generate them by running "python canary/origin_data/inject.py".

– To compute the canary leakage rate, please run "bash canary/eval-leakage-rate/run_pii_detect_step2_infbaseline.sh".

• Hyper-parameter Analysis (Table 5).

– To perform fine-tuning under different hyper-parameter settings, please run "bash scripts-run-finetune/privcode_utility_boosting_hyper.sh". The main adjustable parameters include the maximum lambda λ_{\max} , privacy budget ϵ , and BERTScore threshold τ_s ,

which can be set by modifying “MAX_LAMBDA”, “TARGET_EPSILON”, and “SIM_THRESHOLD”.

D. Acknowledgement

We thank the authors of `fast-differential-privacy`¹⁵ for providing the code for the DP LLM training framework.

E. Copyright

We have migrated the code repository from GitHub to Zenodo. The DOI for the repository is <https://doi.org/10.5281/zenodo.17845610>.

¹⁵<https://github.com/awslabs/fast-differential-privacy>