



Delft University of Technology

Document Version

Final published version

Licence

CC BY

Citation (APA)

Al-Kaswan, A., Spiess, C., Devanbu, P., Van Deursen, A., & Izadi, M. (2026). Model See, Model Do? Exposure-Aware Evaluation of Bug-vs-Fix Preference in Code LLMs. In *Proceedings - 2026 IEEE/ACM 23rd International Conference on Mining Software Repositories, MSR 2026* (pp. 86-97). Association for Computing Machinery (ACM).
<https://doi.org/10.1145/3793302.3793341>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Model See, Model Do? Exposure-Aware Evaluation of Bug-vs-Fix Preference in Code LLMs

Ali Al-Kaswan
Delft University of Technology
Delft, Netherlands
a.al-kaswan@tudelft.nl

Claudio Spiess
University of California, Davis
Davis, USA
cvspiess@ucdavis.edu

Prem Devanbu
University of California, Davis
Davis, USA
ptdevanbu@ucdavis.edu

Arie van Deursen
Delft University of Technology
Delft, Netherlands
arie.vandeursen@tudelft.nl

Maliheh Izadi
Delft University of Technology
Delft, Netherlands
m.izadi@tudelft.nl

Abstract

Large language models are increasingly used for code generation and debugging, but their outputs can still contain bugs that originate from training data. Distinguishing whether an LLM prefers correct code, or a familiar incorrect version might be influenced by what it's been exposed to during training. We introduce an exposure-aware evaluation framework that quantifies how prior exposure to buggy versus fixed code influences a model's preference. Using the ManySStuBs4J benchmark, we apply Data Portraits for membership testing on the Stack-V2 corpus to estimate whether each buggy and fixed variant was seen during training. We then stratify examples by exposure and compare model preference using code completion as well as multiple likelihood-based scoring metrics. We find that most examples (67%) have neither variant in the training data, and when only one is present, fixes are more frequently present than bugs. In model generations, models reproduce buggy lines far more often than fixes, with bug-exposed examples amplifying this tendency and fix-exposed examples showing only marginal improvement. In likelihood scoring, minimum and maximum token-probability metrics consistently prefer the fixed code across all conditions, indicating a stable bias toward correct fixes. In contrast, metrics like the Gini coefficient reverse preference when only the buggy variant was seen. Our results indicate that exposure can skew bug-fix evaluations and highlight the risk that LLMs may propagate memorised errors in practice.

CCS Concepts

• **Software and its engineering** → **Software verification and validation**; • **Computing methodologies** → **Natural language processing**.

Keywords

Large Language Models, bugs, fixes, Memorisation

ACM Reference Format:

Ali Al-Kaswan, Claudio Spiess, Prem Devanbu, Arie van Deursen, and Maliheh Izadi. 2026. Model See, Model Do? Exposure-Aware Evaluation of Bug-vs-Fix Preference in Code LLMs. In *23rd International Conference on Mining Software Repositories (MSR '26)*, April 13–14, 2026, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3793302.3793341>

1 Introduction

Large Language Models (LLMs) have revolutionised code-related tasks in software engineering, serving as powerful tools for code completion, generation, and debugging [15, 17, 18, 51]. Their popularity stems from their ability to understand and produce code in multiple programming languages, often outperforming traditional methods in terms of speed and accessibility [11, 37]. For instance, models like GitHub Copilot, built on top of LLMs, have been widely adopted by developers, with millions of users leveraging them daily for productivity gains [56].

Although LLMs have transformed code generation and debugging workflows, their output is not immune from errors; this introduces substantial risks that can undermine software reliability and security [3, 48]. For example, LLM-generated code has been found to include various kinds of errors [20, 36, 55]. These bugs can propagate through systems, leading to costly failures, data breaches, or exploits in real-world applications. These bugs may stem from coding errors in training data [19, 48].

The ManySStuBs4J dataset collects thousands of SStuBs, ('Simple, Stupid Bugs'), which are single-statement bugs and their fixes from open-source Java repositories [23]. Although seemingly minor, these SStuBs are particularly impactful because they represent real-world errors, mined from open-source projects. Previous work has shown that LLMs are almost twice as likely to produce a buggy version of a line compared to the fixed version [20].

A critical aspect exacerbating the risks associated with LLM-generated code is the issue of memorisation. LLMs can memorise and regurgitate parts of their training data [8, 27, 33, 44]. This issue is particularly pronounced in code LLMs, where training data often consists of vast, publicly sourced, repositories that contain a mix of correct and erroneous code, leading to potential biases in model outputs [3, 19].

For example, if a model has been exposed to a specific buggy code sequence during training, perhaps due to its prevalence in open source datasets, the model may inadvertently favour and propagate code similar to the buggy version, rather than the fixed



This work is licensed under a Creative Commons Attribution 4.0 International License.
MSR '26, Rio de Janeiro, Brazil
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2474-9/26/04
<https://doi.org/10.1145/3793302.3793341>

version [19]; indeed subsequent generations from the model could reflect memorised buggy code rather than correct coding patterns, either learned directly from corrected versions or via appropriate generalizations [3, 4, 49]. If memorised code is regurgitated, this can affect evaluation metrics in benchmarks [3, 49, 53], which raises concerns about the reliability of LLM-assisted code development [4, 53].

Beyond correctness, exposure-influenced behaviour of code LLMs creates concrete security and governance problems. From a security perspective, models that reproduce insecure patterns in training data can introduce systemic vulnerabilities and expand the attack surface across many downstream projects [3, 19, 36, 53]; memorisation of proprietary or sensitive snippets also raises legal risks if such content is regurgitated [3, 4, 25, 26]. Therefore, there is a need for exposure-aware methods that accurately assess how prior exposure to data influences LLMs.

To investigate whether LLMs “see, then do” when it comes to buggy versus fixed code, we develop an exposure-aware evaluation pipeline that samples bug-fix pairs from mined benchmarks, and estimates training exposure for each variant using membership tools. We then measure model preference with a suite of likelihood-based metrics. Additionally we evaluate model completions based on the bug-fix context. We then relate preference to exposure and to bug categories, and perform robustness checks to identify which metrics and categories are most sensitive to exposure, with the goal of disentangling memorisation-driven propagation from genuine learned correctness.

Our evaluation yields several noteworthy findings. First, we find that about 67% of bug-fix pairs in SStuBs were unseen in the model’s training data, and furthermore, the fixes were substantially more likely to have been seen than the bugs. Second, in actual model generations, buggy variants are reproduced far more often than fixes. Moreover, exposure to the buggy variant substantially boosts bug-matching rates in generations, while exposure to the fix provides only marginal improvements in fix-generation rates. These patterns are consistent across evaluated models. Third, likelihood-based metrics such as the minimum and maximum token probability consistently favour the correct, fixed code across all exposure conditions, whereas others like the Gini coefficient can exhibit exposure-driven reversals (preferring the buggy code when the model has seen that bug in training). Finally, model preferences vary by bug category, with some robust to exposure, others highly susceptible to propagation, and a few exhibiting intrinsic bug preference.

We make the following contributions:

Exposure-Aware Evaluation Methodology: We introduce a novel framework to assess code LLM behaviour that accounts for training data exposure, we perform membership inference for each SStuBs sample to determine if the buggy or fixed variant appeared in the model’s training data.

Insights on Metrics and Generations: We compare likelihood-based metrics to identify those sensitive to exposure and break down preferences by bug category.

Empirical Study of Bug Propagation: We conduct an extensive study on the SStuBs benchmark across multiple open-source LLMs, quantifying how prior exposure to buggy or fixed

code influences model preferences and the propensity to generate bugs.

2 Background and Related Work

LLMs have been applied to code completion and repair, but errors in generated code remain a serious concern. Researchers have assembled benchmarks of real bugs to measure repair and detection tools. For example, the ManySStuBs4J dataset collects thousands of single-statement bugs and their fixes from open-source Java repositories [23]. Other standard benchmarks include Defects4J [21] and QuixBugs [28], which capture small but impactful errors (off-by-one mistakes, incorrect conditionals, missing null checks, etc.) in real programs. Recent studies find that large code models still produce buggy code at non-trivial rates: Jesse et al. report that an LLM is almost twice as likely to generate a buggy line (from SStuBs) as the correct fix [20] and that bugs introduced by others take much longer to fix [20]. Security analyses similarly show that model-synthesised code can include insecure patterns such as missing input sanitization or unsafe API usage [36, 48].

A major challenge in evaluating such models is memorisation of training data. Code corpora contain many duplicate snippets, and large models are known to memorise frequent patterns [4, 16, 54]. This raises the classic membership inference problem: given a code snippet, was it in the training set? In practice this is hard to determine, especially in cases where the training data is not available [8, 33, 45].

Prior analyses have shown that training/test overlap can skew evaluation: test snippets discovered in training data can inflate reported performance [24]. For bug-fix pairs, the asymmetry is especially problematic: a bug appears first in a project and may propagate to forks before it is fixed, so the buggy version may be more prevalent in scraped data than the fix. Jahanshahi and Mockus [19] found that 17% of the blobs in the Stack-v2 are copied and have a newer version of the code available. Of these newer versions, around 17% are bug fixes. These ‘orphan vulnerabilities’ can persist in public repositories even after a fix has been issued [19, 40].

Despite these insights, existing work has not explicitly controlled for exposure when comparing models’ behaviour on bugs versus fixes. Previous work notes that models sometimes echo insecure or erroneous patterns from their training data [4, 9, 13, 31, 35, 43, 53]. But none of these quantify how differing exposure to the buggy versus fixed code affects model preferences. In effect, previous evaluations may confound a model’s learned correctness with simple memorisation. In this work, we address these gaps with an exposure-aware analysis of bug-fix preference.

3 Approach

The approach is described in Figure 1. Our exposure-aware evaluation framework studies bug propagation in code LLMs through an end-to-end pipeline: starting from the left, bug-fix pairs are sampled from the SStuBs dataset; exposure is estimated for each variant from the training corpus; samples are stratified into categories; model preferences are evaluated through likelihood-based scoring and actual code generations from LLMs; generations are




Figure 1: Overview of our exposure-aware evaluation framework

matched against the original bug or fix variants; and finally, results are analysed to quantify exposure effects on bug-fix preferences.

To study the impact of exposure on bug propagation, we first sample a number of SStuBs from public repos. We then perform membership inference to check if the buggy or fixed version of the code is present in the training data.

With membership inference we can determine whether a bug or its associated fix appears in the large corpus. The samples can be stratified into four sets based on this analysis:

- **Neither** the bug nor the fix are present in the training data
- **Both** the bug and fix are present¹
- **Only** the **bug** is present, but the fix is not
- **Only** the **fix** is present, but the bug is not

Using a variety of LLMs, we can then calculate the model preference for either the bug or the fix, for each of the previously mentioned sets. We measure both conditional likelihoods and actual generative behaviour. The conditional likelihoods are extracted from the model losses and are scored using different token-level signals. Generation is evaluated using the LLM under decoding.

3.1 Research Questions

To guide our evaluation of code LLMs on bug propagation, we address the following research questions (RQs). These build on our hypothesis that training exposure to buggy or fixed code influences model behaviour.

RQ1: How does prior exposure to buggy or fixed code influence model preferences as measured by likelihood-based metrics? By computing metrics such as token probabilities we investigate biases in model scoring, helping disentangle memorisation-driven preferences from learned generalisations.

RQ2: How do exposure effects on preferences and generations vary by bug category? This RQ breaks down results by SStuBs categories to pinpoint which bug types are most susceptible to memorisation.

RQ3: How does exposure affect the generation of buggy versus fixed code in model completions? By focusing on verbatim regurgitation rates, we quantify explicit bug propagation in sampled outputs, assessing real-world risks like vulnerability introduction in downstream applications.

3.2 Datasets

Our study utilises two primary datasets: Stack-v2 [30] as the training corpus for membership inference and ManySStuBs4J [23] for bug-fix pairs.

To estimate exposure and perform membership inference, we utilise the Stack-v2 dataset. The Stack-v2 is the largest publicly curated code dataset spanning 619 programming languages and totalling 68TB [30]. The Stack-v2 is mined from the Software Heritage archive [1]. Its massive size improves the chance of overlaps with the SStuBs data and makes it suitable for training state-of-the-art LLMs. The Stack-v2 is widely used in research as it supports many languages and excludes restrictively licenced code [30].

SStuBs focus on single-statement bugs that can be fixed with single-line changes. SStuBs allow us to isolate changes without the complexity of multi-line or multi-file changes, which allows us to study model behaviour [21, 23]. We use the ManySStuBs4J dataset, which contains mined bug-fix pairs from open-source Java projects [20, 23]. We opt to use the processed version of the dataset from Jesse et al. as it contains additional information such as the number of commits between the fix and the bug, as well as the location of the bug in the source file [20]. The dataset consists of 16, 899 bug-fix pairs.

3.3 Models

For our evaluation, we focus on models with open weights, and training data where possible, to improve reproducibility. Because of hardware limitations, we omit models that are too expensive to run. Due to hardware limitations we can support models with up to 20B parameters. We found that few state-of-the-art models publish their training data. Of the 50 top performing models on the EvalPlus benchmark [29], only two have open training data, and both are instruction tuned versions of StarCoder2.²

The StarCoder2 family of models is fully open-weight and open-data. The models are based on the Transformer architecture and are available in 3 sizes, 3B, 7B and 15B parameters [30].³

StarCoder2 is trained on the Stack-v2 dataset, which is openly available. The 3B and 7B variants are trained on a subset of the Stack-v2 which only includes 17 programming languages (including Java), while the 15B parameter model is trained on the full set of 600+ languages [30].

¹Note that the number of instances data might differ between the bug and fix

²EvalPlus Benchmark: <https://evalplus.github.io/leaderboard.html> (Accessed October 2025)

³StarCoder2: <https://huggingface.co/collections/bigcode/starcoder2>

We also include Mellum-4B which is an open-weight (but not fully open-data) model trained on Stack-v2 and some other datasets which are not publicly available [38].⁴ Despite its partially undisclosed training data, we can estimate exposure for SStuBs pairs, as the Stack-v2 forms the core of its code training.

Additionally, we include SmolLM3 (3B) [5], a fully open-weight model introduced by Hugging Face in 2025. SmolLM3 was trained on roughly 11 trillion tokens using exclusively public datasets in a staged curriculum that mixes web text, mathematics, and code. The code component of its training is drawn from Stack-v2 and other sources like GitHub issues, notebooks, and Q&A forums [5]. We therefore believe it is unlikely that any of the SStuBs ended up in the training data from sources other than Stack-v2. We use the SmolLM3 base checkpoint, without instruction tuning.⁵

3.4 Data Portraits

To estimate the exposure of code variants, we utilise Data Portraits [32], a membership inference tool designed to measure the presence of specific text snippets or patterns in large corpora. We use the Data Portrait created by the authors for Stack-v2.⁶

Data Portraits use strided Bloom filters [7, 32] to enable efficient membership testing over large datasets. Bloom filters are probabilistic data structures that can quickly determine whether an element is in a set or not using fixed-length hash functions. Because different strings may map to the same position in the filter, Bloom filters can produce false positives, indicating membership when the element is not actually present, but they do not yield false negatives. Consequently, when applied in Data Portraits, some code snippets flagged as members might not truly exist in the corpus [32].

The false positive rate can be reduced by increasing the number of bits output by the hash function. In the case of Data Portrait these are tuned to achieve a false positive rate of 0.1% [32]. This results in a portrait of 90GB (compared to 68TB for the full dataset) for Stack-v2, which fits in memory and allows for almost instant membership testing [32].

It is important to note a key constraint of this method: its soundness is only guaranteed for queries exceeding a specific length. This results from the construction of the Data Portrait [32]. The tool was built using a fixed n-gram *width* (w) of 50 tokens and a *stride* (s) of 50 tokens. This means only 50-token sequences starting at every 50-token interval in the training corpus were hashed into the Bloom filter. Consequently, a snippet could fall entirely between these sampling points, leading to a false negative (i.e., the tool reporting a snippet as ‘unseen’ despite its presence in the data).

To ensure that any query is guaranteed to fully contain at least one of the sampled n-grams, regardless of its alignment, its length (L) must satisfy the condition $L \geq s + w - 1$. For the Stack-v2 Data Portrait, with $s = 50$ and $w = 50$, the minimum required query length is $50 + 50 - 1 = 99$ tokens. Any query of 99 tokens or more is thus guaranteed to contain at least one complete 50-token n-gram that would have been hashed and captured in the Portrait, ensuring the test is sound. We adhere to this constraint for all our membership queries by padding any shorter SStuB to 99 tokens

using the surrounding context. In the case of Data Portraits, every character, except spaces and newlines, is considered a token.

We pad from both sides to ensure that we capture as much of the bug or fix behaviour of the snippet. If a sample is only padded with its preceding context, a hit in the Bloom filter could only cover the context and none of the sample. By padding both sides, we can ensure that the hit covers at least $(99 - sample_len)/2$ tokens of the sample in the worst case.

When a code snippet is queried, the Data Portrait tool calculates a ‘badness’ score. This score represents the fraction of the snippet’s n-grams that are found in the Bloom filter. A higher score indicates a stronger signal that the snippet, or a very similar sequence, is present in the training corpus. For samples with $sample_len \leq 99$ only a single hit is possible.

A high ‘badness’ score can arise from a query accumulating multiple hits. The cause is the duplicated presence of code blocks in the training corpus. For example, a function may appear in multiple files. Because the stride is applied independently to each file, the function’s starting position relative to the beginning of the file dictates which of its internal 50-token n-grams are sampled and hashed. In one file, the stride may capture an n-gram from the beginning of the function, while in another file, a different alignment causes an n-gram from the middle of the same function to be hashed. Consequently, when the full function is used as a query, the sliding-window registers multiple distinct hits against the portrait, one for each independently hashed segment discovered across the corpus [32]. Therefore, a single query can generate numerous hits, indicating that different parts of it were seen by the model due to their varied placement across the training data. This underestimates the number of duplicates in the training data as some duplicates are aligned.

To convert this probabilistic score into a binary seen or unseen classification for our study, we apply a threshold. Following the methodology of Marone and Van Durme, we classify a snippet as exposed (seen) if its badness score is 90% or higher [32].

We apply Data Portraits to determine whether the SStuBs bug–fix pairs from our dataset were present in Stack-v2, the training corpus used for the models in our study [30].

3.5 Model Preference Metrics

To evaluate the preference of code LLMs between buggy and fixed variants, we compute $P(w_t | w_{t-1})$, the conditional probability of each token comprising the bug or fix, given the preceding tokens i.e., the code lines before the buggy or fixed statement. As the preceding context is the same for any bug-fix pair, it follows that the probabilities are equal. Thus, we consider only the conditional probabilities of the tokens *comprising* the bug and fix code, which may have differing lengths. We collect two token probability sequences, one for the bug and one for the fix, per model and SStuBs pair.

We study multiple likelihood-based metrics to determine model preference based off the token probabilities. The model’s preferred variant is selected based on the metric’s orientation. For example, if the mean of the token probabilities comprising the fix is higher than the bug, the model prefers the fix. In contrast, a lower perplexity (surprisal) indicates greater preference over a higher perplexity sequence.

⁴Mellum-4B: <https://huggingface.co/JetBrains/Mellum-4b-base>

⁵SmolLM3-3B: <https://huggingface.co/HuggingFaceTB/SmolLM3-3B-Base>

⁶Stack-v2 Data Portrait: <https://stack-v2.dataportraits.org>

However, the choice of metric is not obvious, nor has a singular standard been established in prior work [46], as metrics capture different aspects of the token probability distribution. Therefore, we evaluate various metrics collected from prior work, which we introduce below.

Length $\uparrow$ The total number of tokens in the code sequence. While not a direct measure of preference, length can influence probability calculations and is used as a baseline metric [47].

Perplexity $\downarrow$ A widely used metric in language modelling that measures how well the model predicts a given sequence [52]. It is the exponentiation of the average negative log-likelihood per token. Lower perplexity indicates better model confidence and preference for the sequence.

Minimum Probability $\uparrow$ The smallest probability assigned to any token in the sequence. This metric highlights the least confident token prediction, with higher minimum probabilities indicating more uniformly confident predictions.

Maximum Probability $\uparrow$ The highest probability among all tokens in the sequence. This reflects the model’s peak confidence at any position in the code snippet, which may indicate familiarity or memorisation.

Gini Coefficient $\downarrow$ A statistical measure of inequality in the distribution of token probabilities, where values closer to zero indicate a more uniform probability distribution, and higher values indicate more skewed confidence towards certain tokens [6].

Geometric Mean $\uparrow$ The n th root of the product of token probabilities (multiplicative average), which penalises sequences with any low-probability tokens more heavily than high ones, emphasising consistently high or low probabilities across tokens.

Arithmetic Mean $\uparrow$ The arithmetic mean of token probabilities across the sequence. This captures the model’s average confidence in the output, can be skewed by outlier high probability tokens. Lower arithmetic mean indicates lower average confidence (higher uncertainty) across the sequence.

3.6 Model Generations

To complement the likelihood-based metrics, we evaluate the models’ actual generative behaviour. For each bug-fix pair in the SStuBs dataset, we provide the model with the preceding code context (i.e., the code lines before the buggy or fixed statement) as a prompt, limited to 2048 tokens. We then generate completions using *temperature*=0.8 and *top_p* = 0.95 following Lozhkov et al. [12, 30]. We generate five independent completions per prompt, to limit the computational cost. Each completion limited to a maximum of 64 tokens, this is in line with literature [20]. We found that almost all SStuBs were under this limit. We then assess whether these generations match the buggy variant, the fixed variant, both (if different generations match each), or neither. Matching is determined by exact string comparison [20].

We compute the following disjoint aggregate statistics stratified by exposure category, which sum to 100%:

Fix without bugs: Proportion of samples where at least one generation matches the fixed variant but none match the bug.




Figure 2: Distribution of exposure of bugs and fixes in logarithmic scale

Bug without fixes: Proportion of samples where at least one generation matches the buggy variant but none match the fix.

Mix bug-fix: Proportion of samples where at least one generation matches the bug and at least one matches the fix.

No bug no fix: Proportion of samples where no generations match either the buggy or the fixed variant.

3.7 Environment and Data Availability

All experiments were run on a server with two Nvidia A40 GPUs with 48GB of VRAM (driver version 535.261.03), 256GB of RAM and an AMD EPYC 9334 CPU with 64 cores. The experiments were run with Docker. A Dockerfile to reproduce the environment is available in the replication package.

A replication package, including all data and code, is available on Zenodo⁷.

4 Exposure Characterization of Dataset

To understand the baseline exposure landscape, we first analysed the distribution of exposure across all SStuBs samples using Data Portraits membership inference on the Stack-v2 training corpus. Figure 2 illustrates the exposure distribution for both buggy and fixed code variants on a logarithmic scale. We observe that there are four major clusters, the first around a fix and bug exposure score of 0, the second around exposure scores of 1, a third around 1 and 0 and a fourth around 0 and 1 respectively.

There are a few samples with a fix exposure score less than 1 and a bug exposure score of 1. A score between 0 and 1 indicates that part of the sequence had a hit in the Data Portrait. But the length of the hit was not sufficient to conclude membership. Note that we only consider samples where the exposure is 90% or more to be members, as per the original work [32].

⁷Replication package: <https://doi.org/10.5281/zenodo.17423424>

Table 1: Count and mean number of commits until fix for each SStuB exposure category

Category	Count	#Commits	%
Neither seen	11286	4735	67%
Both seen	2169	9328	13%
Only Bug	1109	14152	7%
Only Fix	2335	6373	14%
Total	16899	6169	

Some have exposures higher than 1, this can be caused by the bug or fix appearing multiple times in the training data. Some code snippets were duplicated dozens or even hundreds of times in the training data, likely due to repository forks, copy-pasted code, or independently written clones.

As shown in Table 1, the vast majority of samples (11,286 pairs, or 67% of the dataset) had neither the bug nor the fix present in the training data. This finding suggests that most evaluation instances represent genuinely unseen code patterns. Among pairs with at least one variant exposed, fixes were substantially more likely to appear in the training corpus than bugs. Specifically, 2,335 samples (14% of the dataset) contained only the fix in Stack-v2, while only 1,109 samples (7%) contained exclusively the bug. A smaller subset of 2,169 pairs (13%) had fixed and buggy variants present in the training data.

We examine the temporal dynamics of bug-fix pairs by analysing the number of commits between the introduction of a bug and its subsequent fix. As shown in Table 1, the mean number of commits until fix varies substantially between exposure categories. Pairs where only the bug was seen exhibit the longest bug lifetime, with an average of 14,152 commits before correction, nearly double the dataset mean. In contrast, pairs where neither variant was seen have the shortest bug lifetime (4,735 commits on average). Pairs with both variants present show an intermediate lifetime (9,328 commits), while fix-only pairs average 6,373 commits before the fix.

Most of the bug-fix pairs in our data set were not present in the training corpus, and 67% of the samples had neither variant seen. When exposure exists, fixes appear more frequently than bugs, and fixes tend to be longer on average than their buggy counterparts. Bug-only pairs exhibit nearly double the bug lifetime compared to the dataset mean, while neither-seen pairs show the shortest lifetimes.

5 Results

In this section, we present the empirical findings from our evaluation, addressing each research question in turn. For clarity and brevity, our analysis will primarily focus on the results of the StarCoder2-7B model. We found that the core patterns were remarkably consistent across the three StarCoder2 models evaluated. The complete results are available in the replication package.

5.1 RQ1: Exposure Influence on Metrics

In section 3.5, we introduced the metrics used in our study. As there is no known singular, ideal metric, in this section we evaluate

metric suitability for selecting the fix over the bug. We study which metric reliably prefers fixes over bugs, *and* remains stable when controlling for exposure.

Table 2 present model preferences for pairs where exactly one variant (bug XOR fix) was seen during training. We pick an equal number of samples for each exposure setting (1109). This stratification isolates the effect of asymmetric exposure on model behaviour. When only the fix was present in the training data, the minimum probability showed the strongest bias towards correctness in all models, with 68–72% of the pairs favouring the fix between different models. Maximum probability exhibited similar robustness.

When only the bug was seen, minimum probability maintained its preference for fixes (77–79% across models), resisting exposure bias. This suggests that minimum probability captures something fundamental about code correctness beyond simple pattern matching. Maximum probability also remained fix-preferring in bug-only scenarios (60–70%), though with slightly reduced strength compared to fix-only cases.

In contrast, Gini coefficient and arithmetic mean showed exposure-driven preference reversal. When only the fix was seen, these metrics moderately preferred fixes (60–62%), but when only the bug was seen, they flipped to preferring bugs (56–62%). Length showed minimal discriminative power in either direction, hovering near 50–60% across all conditions.

Perplexity reveals varying preferences. In fix-only conditions, all models favoured fixes. Under bug-only exposure, StarCoder2-7B and SmolLM3 shifted to prefer bugs, while Mellum-4B maintained a fix preference.

To further isolate learned correctness from exposure, Table 3 present results for pairs where both variants were seen (AND condition) or neither was seen (NOR condition).

Across all three models, minimum probability maintained strong fix preference in both symmetric conditions (73–79% for both-seen, 76–79% for neither-seen). Gini coefficient and arithmetic mean lost nearly all discriminative power in symmetric conditions, hovering near or below 50% fix preference. Geometric mean occupied a middle ground, showing moderate fix preference in neither-seen conditions but losing discriminative power in both-seen scenarios. Perplexity preferences were more balanced or fix-leaning. Both-seen scenarios show near parity. Neither-seen conditions consistently preferred fixes across models.

The preference patterns observed for StarCoder2-7B (Table 2) replicate almost identically in Mellum-4B (Table 3), despite Mellum being trained on additional proprietary data beyond Stack-v2. SmolLM3, though showing slightly weaker overall fix preference, exhibits the same relative ordering of metrics and the same exposure-driven reversals for Gini and arithmetic mean.

Minimum probability and maximum probability consistently prefer fixes over bugs across all exposure conditions. In contrast, metrics like Gini coefficient and arithmetic mean show strong susceptibility to exposure bias, reversing their preference toward bugs when only the buggy variant is seen.

5.2 RQ2: Exposure Effects by Bug Category


To investigate how bug-vs-fix preference varies by bug type, we stratified our analysis by the SStuBs bug categories for pairs where

Table 2: Proportion of Preferred Variants by Bug xor Fix Exposure Across Models


Model	StarCoder2-7B				Mellum-4B				SmolLM3			
Exposure	Fix		Bug		Fix		Bug		Fix		Bug	
Preferred	Fix	Bug	Fix	Bug	Fix	Bug	Fix	Bug	Fix	Bug	Fix	Bug
length	0.59	0.41	0.48	0.52	0.59	0.41	0.48	0.52	0.52	0.48	0.49	0.51
perplexity	0.62	0.38	0.44	0.56	0.53	0.47	0.58	0.42	0.54	0.46	0.44	0.56
min_prob	0.68	0.32	0.78	0.22	0.70	0.30	0.77	0.23	0.72	0.28	0.79	0.21
max_prob	0.58	0.42	0.70	0.30	0.57	0.43	0.69	0.31	0.54	0.46	0.60	0.40
gini	0.60	0.40	0.40	0.60	0.60	0.40	0.41	0.59	0.49	0.51	0.39	0.61
geometric_mean	0.62	0.38	0.44	0.56	0.64	0.36	0.47	0.53	0.54	0.46	0.44	0.56
arithmetic_mean	0.60	0.40	0.38	0.62	0.62	0.38	0.39	0.61	0.50	0.50	0.40	0.60

Table 3: Proportion of Preferred Variants by Both and Neither Exposure Across Models

Model	StarCoder2-7B				Mellum-4B				SmolLM3			
Exposure	Both		Neither		Both		Neither		Both		Neither	
Preferred	Fix	Bug	Fix	Bug	Fix	Bug	Fix	Bug	Fix	Bug	Fix	Bug
length	0.53	0.47	0.54	0.46	0.53	0.47	0.54	0.46	0.52	0.48	0.53	0.47
perplexity	0.51	0.49	0.57	0.43	0.53	0.47	0.58	0.42	0.49	0.51	0.56	0.44
min_prob	0.73	0.27	0.76	0.24	0.73	0.27	0.76	0.24	0.77	0.23	0.79	0.21
max_prob	0.65	0.35	0.63	0.37	0.65	0.35	0.63	0.37	0.56	0.44	0.62	0.38
gini	0.48	0.52	0.54	0.46	0.48	0.52	0.54	0.46	0.46	0.54	0.52	0.48
geometric_mean	0.51	0.49	0.57	0.43	0.51	0.49	0.57	0.43	0.49	0.51	0.56	0.44
arithmetic_mean	0.49	0.51	0.54	0.46	0.49	0.51	0.54	0.46	0.46	0.54	0.52	0.48

**Figure 3: Model preference per bug category (StarCoder2-7B)**

either the bug or the fix was seen. Figure 3, 4, and 5 show the preference breakdown for StarCoder2-7B, Mellum-4B, and SmolLM3, respectively. Each figure contrasts a metric which is biased towards fixes (min prob, top row) with an exposure-sensitive metric (geometric mean, bottom row) under two asymmetric exposure conditions: ‘Seen: fix’ (left) and ‘Seen: bug’ (right), note that these samples are exclusive, they belong to only one of the two exposure criteria, similar to Table 2. The blue bars represent the proportion of pairs where the fix was preferred, while the red bars indicate a preference for the bug.

**Figure 4: Model preference per bug category (Mellum-4B)**

Across all models, the results confirm the findings from RQ1 at a more granular level: min prob consistently favours fixes even when only the bug was seen, whereas geometric mean often flips its preference to align with the exposed variant. For instance, in Figure 3, when only the bug is seen (top right), minprob still prefers the fix in most categories, such as DIFFERENT METHOD SAME ARGS (188 fix vs. 64 bug). In contrast, geometric mean’s preference for many of these same categories shifts dramatically towards the bug (bottom right), as seen with CHANGE IDENTIFIER (94 fix vs. 122 bug).




Figure 5: Model preference per bug category (SmolLM3)

Certain bug categories demonstrate strong robustness to exposure. For CHANGE OPERATOR, all models show a strong and stable preference for the fix, regardless of which variant was seen or which metric was used.

Conversely, some categories are susceptible to exposure-driven propagation. The CHANGENUMERAL category is a prime example. For all three models, both metrics show a preference flip. With Mellum-4B (Figure 4), min prob preference shifts from 27-8 in favour of the fix to 11-10. The effect is even more pronounced for geometric mean, which shifts from 87-31 (fix) to 56-102 (bug). Similar flips are observed for LESS SPECIFIC IF and OVERLOAD METHOD MORE ARGS.

Finally, we identify categories where the model consistently prefers the bug, regardless of exposure. For CHANGE UNARY OPERATOR, all models show a strong preference for the buggy variant even when only the fix was seen. For Mellum-4B, the preference is 5-18 in favour of the bug. This suggests the ‘buggy’ pattern may be far more common in the general training data or represents a strong prior that limited exposure to the fix cannot overcome. These patterns of category-specific robustness and susceptibility are consistent across all three models.

Model preference varies significantly by bug category. Some categories are robust, with models consistently preferring the fix regardless of exposure. Others are more sensitive to exposure. A few categories, show an intrinsic model preference for the buggy version, which persists even when only the fix is seen.

5.3 RQ3: Exposure Impact on Generations

While RQ1 examines model preferences through likelihood-based metrics, RQ3 investigates actual generative behaviour by analysing what code models produce. We evaluate whether completions match the buggy variant, the fixed variant, both, or neither, stratifying results by exposure category to measure memorisation effects on generation.

The results for StarCoder2-7B (we omit Mellum-4B as the results were similar to the StarCoder2 models) and SmolLM3 are presented in Figure 6 and 7.




Figure 6: Model generations per exposure category (StarCoder2-7B)




Figure 7: Model generations per exposure category (SmolLM3)

Across all models and exposure categories, generations match buggy code at substantially higher rates than fixed code, with bug-exclusive matches consistently outpacing fix-exclusive matches. For StarCoder2-7B (Figure 6), ‘Bug without fixes’ rates were highest in bug-only exposure and lowest in neither-seen, while ‘Fix without bugs’ rates were substantially lower across all categories. Mixed matches were moderate, and no-match rates were highest in the neither-seen category.

SmolLM3 (Figure 7) shows lower overall match rates compared to StarCoder2, with similar patterns: ‘Bug without fixes’ highest in bug-only and lowest in neither-seen, ‘Fix without bugs’ consistently low, mixed matches minimal, and no-match rates predominant, especially in neither-seen.

The bug-only exposure category consistently showed the highest ‘Bug without fixes’ rates across models. Conversely, fix-only exposure leads to modest increases in ‘Fix without bugs’ for StarCoder2,

but minimal or no increase for SmolLM3. The both-seen category exhibited high ‘Bug without fixes’ rates for StarCoder2, comparable to bug-only, while SmolLM3’s rates fall between bug-only and neither-seen.

Models generate buggy variants at substantially higher rates than fixes across all exposure categories, with bug-exclusive matches consistently outpacing fix-exclusive ones. Bug-only exposure amplifies bug propagation the most, while neither-seen minimises it; fix-only provides minimal improvement in fix generation.

6 Discussion

Our investigation reveals that code language models, despite encountering fewer bugs than fixes in their training data overall, exhibit a pronounced tendency to propagate bugs over fixes, heavily influenced by specific exposure patterns. Across all models, buggy variants were generated at substantially higher rates than fixes in every exposure category, with bug-only exposure exacerbating this bias. Addressing the core question of if “models see, then do”: models only adhere to this principle to a certain extent. Exposure to bugs strongly drives increased bug generation and preference, whereas fix exposure only weakly promotes corrections. Metrics like minimum probability highlighted that models can generalise toward fixes in all exposure conditions.

The category-specific susceptibilities observed in RQ2, suggest that certain bug types align more closely with prevalent coding idioms or anti-patterns embedded in the training data. For instance, numeral changes may resemble common refactoring or optimisations that appear frequently in code repositories, making models more prone to propagating buggy variants when they are the only ones seen. This alignment could amplify security risks in LLM-generated code, as models might inadvertently reproduce vulnerable patterns. To mitigate this, targeted fine-tuning on under-represented fix patterns could bolster model robustness.

Another insight emerges from contrasting likelihood-based evaluation with actual generative behaviour. We observed that metrics like minimum token probability consistently favour the fixed code across all exposure conditions, whereas in practice the models often generate the buggy variant at much higher rates, especially under bug-only exposure. This reflects that these methods probe different aspects of the model. Likelihood metrics examine the probability distribution the model has learned, whereas generation metrics reflect the model’s actual output under a given decoding scheme. In other words, the model “knows” that the fix should be more probable, but the process of sampling and decoding can still yield the bug as output. Both evaluation approaches can be complementary: likelihood-based metrics can diagnose what the model has learned and its confidence profile, whereas generation-based metrics reveal its practical behaviours and risks.

Bug Lifecycle. Our results reveal that fixes appear more frequently in the training corpus than bugs, with 14% of pairs having only the fix exposed, compared to 7% exposing only the bug. This asymmetry may stem from the inherent construction of datasets like SStuBs, where bugs are mined from commit histories and can

only be identified retrospectively once they are fixed. Bugs, by nature, have a limited observable lifetime in repositories, they persist until corrected, after which the fixed version dominates subsequent versions and forks. Consequently, training corpora like Stack-v2 are likely to capture more instances of the post-fix code, especially in actively maintained projects where fixes propagate quickly through merges and updates.

Analysis of bug lifetimes reveals that detected bugs in exposure categories like bug-only pairs persist for substantially longer compared to fixes or neither-seen pairs. Cases where both variants are seen fall in between, suggesting a detection bias toward longer-lived bugs that survive multiple repository iterations before correction. Very short-lived bugs, which are quickly fixed within a few commits, may evade mining in datasets like SStuBs altogether. This could imply that our evaluation under-represents transient errors, where bugs have comparably short lifetimes, potentially skewing exposure-aware assessments toward more entrenched, harder-to-fix patterns that models are more likely to propagate.

The observed preference for bugs in model generations (RQ3), even in scenarios where fixes are more exposed. Fix-only pairs show only marginal increases in fix-matching rates, this highlights a potential imbalance in training data where erroneous patterns are reinforced through sheer volume or recency bias in code repositories. Bugs, often introduced in recent commits before being fixed, may appear in more transient or duplicated contexts across forks and pull requests. In contrast, fixes, while more stable, might be diluted by the nature of LLM training, where high-frequency error patterns overshadow corrections.

Higher bug generation rates may also arise from biases in dataset mining and the inherent ‘naturalness’ of human-written errors. SStuBs mining favours detectable bugs that were eventually fixed, potentially underrepresenting fixes that resemble common idioms, leading to an evaluation set where bugs appear more generation-prone. Additionally, these bugs were originally authored by programmers, making them feel ‘natural’ in probabilistic terms; models may favour them as they align with frequent coding slips observed in training. Furthermore, these are coding slips that managed to get past unit-testing, and perhaps code review, into the code base. On the other hand, their eventual discovery and correction indicates they were detected, and fixed.

Models. Interestingly, our analysis uncovered minimal difference between the 3B and 7B and 15B variants of StarCoder2, despite the substantial disparity in model size and capacity. The models exhibit comparable generation rates favouring bugs across exposure categories. Preference metrics like minimum probability further reinforced this uniformity. This lack of differentiation suggests that scaling model parameters alone may not suffice to overcome inherent biases in training data.

The near-identical performance of Mellum-4B and the StarCoder2 models across RQ1–RQ3, in contrast to the distinct patterns in SmolLM3, suggests that, for our evaluation, additional training data has minimal impact. Mellum-4B, while incorporating some closed-source datasets beyond Stack-v2, exhibits preference metrics and generation rates that mirror StarCoder2-7B closely, implying that these proprietary additions are limited in scope or relevance to code tasks. SmolLM3, however, deviates with lower overall match

rates and weaker fix preferences, likely due to its significant inclusion of non-code training data; although code is not expected in these portions, incidental inclusions or domain shifts could dilute its code-specific capabilities.

From an AI safety perspective, the tendency of models to regenerate seen bugs, as evidenced by higher bug-matching rates in generations (RQ3), raises significant concerns about downstream harms, such as introducing vulnerabilities into production software. When models memorise and reproduce buggy patterns from training data, they risk propagating errors across multiple projects. This phenomenon aligns with broader discussions on memorisation risks in LLMs, where unintended data regurgitation can undermine software reliability [4, 8, 19].

While developers typically have little control over the composition of pre-training datasets, our study highlights the specific risks that models may inherit from such data. Understanding these exposure effects is therefore essential, particularly because practitioners often deploy models without direct visibility into or control over their training sources [2, 3, 50].

6.1 Threats to Validity

Although we have made efforts to ensure the robustness and reliability of our study, several potential threats to validity should be acknowledged.

Internal Validity. One threat is the potential for false positives in membership inference using Data Portraits, as Bloom filters can incorrectly indicate exposure due to hash collisions, with a reported false positive rate of 0.1% [32]. This could lead to misclassification of bug-fix pairs into exposure categories, potentially exaggerating exposure-driven effects. We limited this impact by adhering to the tool's recommended query length threshold (99 tokens) and using the recommended 90% badness score cutoff for binary classification [32].

A potential limitation in our analysis stems from the context provided to models, which may not be sufficiently constraining to unambiguously distinguish buggy from fixed code. We evaluate single-statement changes with the preceding file context, without full project context, meaning that the preceding lines alone might admit multiple valid completions. This could explain instances where generations match neither variant (prevalent in RQ3), as models generate plausible but unrelated code.

Another threat arises from the probabilistic nature of LLM generations in RQ3, where decoding may introduce variability due to model stochasticity. We mitigated this by generating five independent completions per prompt and aggregating statistics ('Any bug-fix' rates), reducing the influence of single-run anomalies.

External Validity. We rely on Java-specific bug-fix pairs from the SStuBs dataset, which may not generalise to other programming languages with different syntax or error patterns (such as dynamic languages like Python). This could limit the applicability of category-specific findings in RQ2 to multilingual settings. We limited this impact by selecting SStuBs for its focus on single-statement bugs, which are common across languages, and by using models trained on Stack-v2's diverse corpus (619 languages). However, we accepted the Java focus as a necessary scope limitation,

given SStuBs' established use in code LLM evaluations [20, 23, 41], and recommend future extensions to multilingual benchmarks.

We only evaluated a small set of open-weight models (StarCoder2 variants, Mellum-4B, SmolLM3), which may not represent closed-source models like GitHub Copilot or larger architectures. Differences in training data or scale could alter exposure effects. We mitigated this by including models of varying sizes (3B to 15B parameters). We did not evaluate the models using the fill-in-the-middle (FIM) objective. SmolLM3 does not support FIM [5], and StarCoder2-15B exhibits poor FIM performance due to known training errors by its authors [30].

Construct Validity. Our binary classification of exposure based on Data Portraits oversimplifies memorisation, ignoring nuances like partial matches or contextual variations in training data. This could misrepresent true exposure, affecting downstream analyses. We accepted some simplification as a trade-off for scalability, given the size of Stack-v2.

Our analysis measures memorization using literal string matching, as partial or approximate matching lacks a well-defined and reliable criterion for automated data extraction [14]. Because buggy and fixed versions of code often differ by only a small number of tokens, generations that do not exactly match either version are difficult to classify, potentially leading to underestimation of partial memorization. We apply this strict criterion uniformly across all exposure categories, preserving the validity of relative comparisons.

6.2 Future Work

Our exposure-aware evaluation framework opens several avenues for future research.

One promising direction is extending the analysis to multilingual and multi-dataset settings to enhance external validity. While our study focused on Java bug-fix pairs from SStuBs, incorporating benchmarks like Defects4J [21], QuixBugs [28, 39] or PySStuBS [22], which span languages like Python and C++, could reveal language-specific exposure effects. For instance, dynamic languages might exhibit different bug lifetimes or category susceptibilities (such as type-related errors in RQ3), potentially modulated by syntactic differences. By applying Data Portraits to the full Stack-v2 corpus, future work could quantify cross-language propagation risks, providing a more comprehensive view of how code LLMs handle diverse error patterns.

Improving membership inference techniques represents another key area, and would solve the internal validity threats from false positives and binary classifications in Data Portraits. Advanced methods, such as embedding-based similarity searches or neural membership inference [16, 34, 42, 45], could offer finer-grained exposure estimates, accounting for partial matches or contextual variations in training data. This would enable nuanced stratifications beyond our seen/unseen dichotomy, such as weighting by exposure frequency or recency, and better isolate memorisation from generalisation in preference metrics. Black-box membership inference could allow for the detection of exposure, purely based on model characteristics. However, current model-only membership inference approaches are not reliable enough to assist in our evaluation [10, 33].

Finally, the metrics discussed in RQ1 and RQ2 might offer practical utility for developers by helping rank and filter multiple code completion candidates to select the most reliable option. Furthermore, metrics like the Gini coefficient can function as a form of membership inference, allowing developers to identify, and decide what to do with, code that has been memorised directly from the training data, without access to the training corpus.

7 Conclusion

In this paper, we explored whether code LLMs “see, then do” by propagating bugs from training data, hypothesising that exposure to buggy versus fixed code could drive memorisation and error regurgitation. Through an exposure-aware pipeline combining membership inference, likelihood metrics, and generation analysis on the SStuBs benchmark, we empirically examined this across multiple models.

Our results confirm that models do “see, then do” bugs: while many bug-fix pairs remain unseen and fixes are generally more exposed than bugs, models often prefer and generate buggy variants, especially when exposed to them. Some likelihood preferences lean toward fixes but can reverse under bug exposure, with effects varying by bug type.

Our study underscores the need for exposure-aware evaluation in coding LLMs. The “see, then do” effect reveals that models internalise fixes as preferable in likelihood scores, yet frequently reproduce memorised bugs in generations. This discrepancy between knowledge and outputs has implications for tool reliability, leading us to advocate membership testing and diverse metrics to uncover memorisation biases. By controlling for exposure, researchers can better assess true generalisation, mitigate error propagation, and promote safer software development.

References

- [1] Jean-François Abramatic, Roberto Di Cosmo, and Stefano Zacchiroli. 2018. Building the universal archive of source code. *Commun. ACM* 61, 10 (2018), 29–31.
- [2] Ali Al-Kaswan, Sebastian Deatc, Begüm Koç, Arie van Deursen, and Maliheh Izadi. 2025. Code Red! On the Harmfulness of Applying Off-the-shelf Large Language Models to Programming Tasks. *Proceedings of the ACM on Software Engineering* 2, FSE (2025), 2477–2499.
- [3] Ali Al-Kaswan and Maliheh Izadi. 2023. The (ab)use of Open Source Code to Train Large Language Models. In *Proceedings of the 2nd International Workshop on Natural Language-based Software Engineering (NLBSE)*.
- [4] Ali Al-Kaswan, Maliheh Izadi, and Arie Van Deursen. 2024. Traces of memorisation in large language models for code. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, 1–12.
- [5] Elie Bakouch, Loubna Ben Allal, Anton Lozhkov, Nouamane Tazi, Lewis Tunstall, Carlos Miguel Patino, Edward Beeching, Aymeric Roucher, Aksel Joonas Reedi, Quentin Gallouédec, Kashif Rasul, Nathan Habib, Clémentine Fourrier, Hynek Kydlíček, Guilherme Penedo, Hugo Larcher, Mathieu Morlon, Vaibhav Srivastav, Joshua Lochner, Xuan-Son Nguyen, Colin Raffel, Leandro von Werra, and Thomas Wolf. 2025. SmolLM3: smol, multilingual, long-context reasoner. <https://huggingface.co/blog/smollm3>.
- [6] Maciej Besta, Julia Barth, Eric Schreiber, Ales Kubicek, Afonso Catarino, Robert Gerstenberger, Piotr Nyczyk, Patrick Iff, Yueling Li, Sam Houliston, Tomasz Sternal, Marcin Copik, Grzegorz Kwaśniewski, Jürgen Müller, Łukasz Flis, Hannes Eberhard, Hubert Niewiadomski, and Torsten Hoefer. 2025. Reasoning Language Models: A Blueprint. doi:10.48550/arXiv.2501.11223 arXiv:2501.11223 [cs]
- [7] Burton H Bloom. 1970. Space/time trade-offs in hash coding with allowable errors. *Commun. ACM* 13, 7 (1970), 422–426.
- [8] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. 2021. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, 2633–2650.
- [9] Dario Di Palma, Felice Antonio Merra, Maurizio Sfili, Vito Walter Anelli, Fedelucio Narducci, and Tommaso Di Noia. 2025. Do llms memorize recommendation datasets? a preliminary study on movielens-1m. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2582–2586.
- [10] Jan Dubiński, Antoni Kowalczyk, Stanisław Pawlak, Przemysław Rokita, Tomasz Trzciński, and Paweł Morawiecki. 2023. Towards More Realistic Membership Inference Attacks on Large Diffusion Models. *arXiv preprint arXiv:2306.12983* (2023).
- [11] Sarah Fakhoury, Aaditya Naik, Georgios Sakkas, Saikat Chakraborty, and Shuvendu K. Lahiri. 2024. LLM-Based Test-Driven Interactive Code Generation: User Study and Empirical Evaluation. *IEEE Transactions on Software Engineering* 50, 9 (2024), 2254–2268. doi:10.1109/TSE.2024.3428972
- [12] Alex Gu, Baptiste Rozière, Hugh Leather, Armando Solar-Lezama, Gabriel Synnaeve, and Sida I Wang. 2024. Cruxeval: A benchmark for code reasoning, understanding and execution. *arXiv preprint arXiv:2401.03065* (2024).
- [13] Valentin Hartmann, Anshuman Suri, Vincent Bindschaedler, David Evans, Shruti Tople, and Robert West. 2023. SoK: Memorization in General-Purpose Large Language Models. arXiv:2310.18362 [cs.CL] <https://arxiv.org/abs/2310.18362>
- [14] Jamie Hayes, Marika Swanberg, Harsh Chaudhari, Itay Yona, Iliia Shumailov, Milad Nasr, Christopher A Choquette-Choo, Katherine Lee, and A Feder Cooper. 2025. Measuring memorization in language models via probabilistic extraction. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 9266–9291.
- [15] Xinyi Hou, Yanjie Zhao, Yue Liu, Zhou Yang, Kailong Wang, Li Li, Xiapu Luo, David Lo, John Grundy, and Haoyu Wang. 2023. Large language models for software engineering: A systematic literature review. *arXiv preprint arXiv:2308.10620* (2023).
- [16] Hongsheng Hu, Zoran Salicic, Lichao Sun, Gillian Dobbie, Philip S Yu, and Xuyun Zhang. 2022. Membership inference attacks on machine learning: A survey. *ACM Computing Surveys (CSUR)* 54, 11s (2022), 1–37.
- [17] Maliheh Izadi, Roberta Gismondi, and Georgios Gousios. 2022. Codefill: Multi-token code completion by jointly learning from structure and naming sequences. In *Proceedings of the 44th international conference on software engineering*, 401–412.
- [18] Maliheh Izadi, Jonathan Katzy, Tim Van Dam, Marc Otten, Razvan Mihai Popescu, and Arie Van Deursen. 2024. Language models for code completion: A practical evaluation. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, 1–13.
- [19] Mahmoud Jahanshahi and Audris Mockus. 2025. Cracks in the stack: Hidden vulnerabilities and licensing risks in llm pre-training datasets. In *2025 IEEE/ACM International Workshop on Large Language Models for Code (LLM4Code)*. IEEE, 104–111.

- [20] Kevin Jesse, Toufique Ahmed, Premkumar T Devanbu, and Emily Morgan. 2023. Large language models and simple, stupid bugs. In *2023 IEEE/ACM 20th International Conference on Mining Software Repositories (MSR)*. IEEE, 563–575.
- [21] René Just, Dariosush Jalali, and Michael D Ernst. 2014. Defects4J: A database of existing faults to enable controlled testing studies for Java programs. In *Proceedings of the 2014 international symposium on software testing and analysis*. 437–440.
- [22] Arthur V Kamiński, Luisa Palechor, Cor-Paul Bezemer, and Abram Hindle. 2021. Pysstubs: Characterizing single-statement bugs in popular open-source python projects. In *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. IEEE, 520–524.
- [23] Rafael-Michael Karampatsis and Charles Sutton. 2020. How often do single-statement bugs occur? the manystubs4j dataset. In *Proceedings of the 17th international conference on mining software repositories*. 573–577.
- [24] Anjan Karmakar, Julian Aron Prenner, Marco D'Ambros, and Romain Robbes. 2022. Codex hacks hackerrank: Memorization issues and a framework for code synthesis evaluation. *arXiv preprint arXiv:2212.02684* (2022).
- [25] Jonathan Katzy, Razvan Popescu, Arie Van Deursen, and Maliheh Izadi. 2024. An Exploratory Investigation into Code License Infringements in Large Language Model Training Datasets. In *Proceedings of the 2024 IEEE/ACM First International Conference on AI Foundation Models and Software Engineering*. 74–85.
- [26] Jonathan Katzy, Razvan Mihai Popescu, Arie van Deursen, and Maliheh Izadi. 2025. The Heap: A Contamination-Free Multilingual Code Dataset for Evaluating Large Language Models. In *Proceedings of the 2025 IEEE/ACM Second International Conference on AI Foundation Models and Software Engineering*. <https://arxiv.org/abs/2501.09653>
- [27] Jiaolong Kong, Xiaofei Xie, and Shangqing Liu. 2025. Demystifying Memorization in LLM-Based Program Repair via a General Hypothesis Testing Framework. *Proceedings of the ACM on Software Engineering* 2, FSE (2025), 2712–2734.
- [28] Derrick Lin, James Koppel, Angela Chen, and Armando Solar-Lezama. 2017. QuixBugs: A multi-lingual program repair benchmark set based on the Quixey Challenge. In *Proceedings Companion of the 2017 ACM SIGPLAN international conference on systems, programming, languages, and applications: software for humanity*. 55–56.
- [29] Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2023. Is Your Code Generated by ChatGPT Really Correct? Rigorous Evaluation of Large Language Models for Code Generation. In *Thirty-seventh Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=1qv6x10Cu7>
- [30] Anton Lozhkov, Raymond Li, Loubna Ben Allal, Federico Cassano, Joel Lamy-Poirier, Nouamane Tazi, Ao Tang, Dmytro Pykhtar, Jiawei Liu, Yuxiang Wei, Tianyang Liu, Max Tian, Denis Kocetkov, Arthur Zucker, Younes Belkada, Zijian Wang, Qian Liu, Dmitry Abulkhanov, Indraneil Paul, Zhuang Li, Wen-Ding Li, Megan Risdal, Jia Li, Jian Zhu, Terry Yue Zhuo, Evgenii Zheltonozhskii, Nii Osae Osae Dade, Wenhao Yu, Lucas Krauß, Naman Jain, Yixuan Su, Xuanli He, Manan Dey, Edoardo Abati, Yekun Chai, Niklas Muennighoff, Xiangru Tang, Muh-tasham Oblokulov, Christopher Akiki, Marc Marone, Chenghao Mou, Mayank Mishra, Alex Gu, Binyuan Hui, Tri Dao, Armel Zebaze, Olivier Dehaene, Nicolas Patry, Canwen Xu, Julian McAuley, Han Hu, Torsten Scholak, Sébastien Paquet, Jennifer Robinson, Carolyn Jane Anderson, Nicolas Chapados, Mostofa Patwary, Nima Tajbakhsh, Yacine Jernite, Carlos Muñoz Ferrandis, Lingming Zhang, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. 2024. StarCoder 2 and The Stack v2: The Next Generation. *arXiv:2402.19173 [cs.SE]* <https://arxiv.org/abs/2402.19173>
- [31] Xingjun Ma, Yifeng Gao, Yixu Wang, Ruofan Wang, Xin Wang, Ye Sun, Yifan Ding, Hengyuan Xu, Yunhao Chen, Yunhan Zhao, et al. 2025. Safety at Scale: A Comprehensive Survey of Large Model Safety. *CoRR* (2025).
- [32] Marc Marone and Benjamin Van Durme. 2023. Data portraits: Recording foundation model training data. *Advances in Neural Information Processing Systems* 36 (2023), 15121–15135.
- [33] Matthieu Meeus, Igor Shilov, Shubham Jain, Manuel Faysse, Marek Rei, and Yves-Alexandre de Montjoye. 2025. Sok: Membership inference attacks on llms are rushing nowhere (and how to fix it). In *2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*. IEEE, 385–401.
- [34] Fatemehsadat Miresghallah, Kartik Goyal, Archit Uniyal, Taylor Berg-Kirkpatrick, and Reza Shokri. 2022. Quantifying Privacy Risks of Masked Language Models Using Membership Inference Attacks. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 8332–8347. <https://aclanthology.org/2022.emnlp-main.570>
- [35] Fatemehsadat Miresghallah, Archit Uniyal, Tianhao Wang, David Evans, and Taylor Berg-Kirkpatrick. 2022. Memorization in NLP Fine-tuning Methods. In *First Workshop on Pre-training: Perspectives, Pitfalls, and Paths Forward at ICML 2022*.
- [36] Ahmad Mohsin, Helge Janicke, Adrian Wood, Iqbal H Sarker, Leandros Maglaras, and Naeem Janjua. 2024. Can we trust large language models generated code? a framework for in-context learning, security patterns, and code evaluations across diverse llms. *arXiv preprint arXiv:2406.12513* (2024).
- [37] Daye Nam, Andrew Macvean, Vincent Hellendoorn, Bogdan Vasilescu, and Brad Myers. 2024. Using an LLM to Help With Code Understanding. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering (Lisbon, Portugal) (ICSE '24)*. Association for Computing Machinery, New York, NY, USA, Article 97, 13 pages. doi:10.1145/3597503.3639187
- [38] Nikita Pavlichenko, Iurii Nazarov, Ivan Dolgov, Ekaterina Garanina, Karol Lasocki, Julia Reshetnikova, Sergei Boitsov, Ivan Bondyrev, Dariia Karaeva, Maksim Sheptyakov, Dmitry Ustalov, Artem Mukhin, Semyon Proshev, Nikita Abramov, Olga Kolomytseva, Kseniia Lysaniuk, Ilia Zavidnyi, Anton Semenkina, Vladislav Tankov, and Uladzislau Sazanovich. 2025. Mellum-4b-base.
- [39] Julian Aron Prenner, Hlib Babii, and Romain Robbes. 2022. Can OpenAI's codex fix bugs? an evaluation on QuixBugs. In *Proceedings of the Third International Workshop on Automated Program Repair*. 69–75.
- [40] David Reid, Mahmoud Jahanshahi, and Audris Mockus. 2022. The extent of orphan vulnerabilities from code reuse in open source software. In *Proceedings of the 44th international conference on software engineering*. 2104–2115.
- [41] Cedric Richter and Heike Wehrheim. 2022. TSSB-3M: Mining single statement bugs at massive scale. In *Proceedings of the 19th International Conference on Mining Software Repositories*. 418–422.
- [42] Ahmed Salem, Yang Zhang, Mathias Humbert, Pascal Berrang, Mario Fritz, and Michael Backes. 2018. ML-Leaks: Model and Data Independent Membership Inference Attacks and Defenses on Machine Learning Models. doi:10.48550/arXiv.1806.01246 *arXiv:1806.01246 [cs]*.
- [43] Fabio Salerno, Ali Al-Kaswan, and Maliheh Izadi. 2025. How much do code language models remember? an investigation on data extraction attacks before and after fine-tuning. In *2025 IEEE/ACM 22nd International Conference on Mining Software Repositories (MSR)*. IEEE, 465–477.
- [44] Avi Schwarzschild, Zhili Feng, Pratyush Maini, Zachary Lipton, and J Zico Kolter. 2024. Rethinking llm memorization through the lens of adversarial compression. *Advances in Neural Information Processing Systems* 37 (2024), 56244–56267.
- [45] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*. IEEE, 3–18.
- [46] Ola Shorinwa, Zhiting Mei, Justin Lidard, Allen Z. Ren, and Anirudha Majumdar. 2025. A Survey on Uncertainty Quantification of Large Language Models: Taxonomy, Open Research Challenges, and Future Directions. *ACM Comput. Surv.* (June 2025). doi:10.1145/3744238
- [47] Claudio Spiess, David Gros, Kunal Suresh Pai, Michael Pradel, Md Rafiqul Islam Rabin, Amin Alipour, Susmit Jha, Prem Devanbu, and Toufique Ahmed. 2024. Calibration and correctness of language models for code. *arXiv preprint arXiv:2402.02047* (2024).
- [48] Florian Tamboon, Arghavan Moradi-Dakheel, Amin Nikanjam, Foutse Khomh, Michel C Desmarais, and Giuliano Antoniol. 2025. Bugs in large language models generated code: An empirical study. *Empirical Software Engineering* 30, 3 (2025), 65.
- [49] Yao Wan, Guanghua Wan, Shijie Zhang, Hongyu Zhang, Yulei Sui, Pan Zhou, Hai Jin, and Lichao Sun. 2024. Does Your Neural Code Completion Model Use My Code? A Membership Inference Approach. *arXiv preprint arXiv:2404.14296* (2024).
- [50] Kun Wang, Guibin Zhang, Zhenhong Zhou, Jiahao Wu, Miao Yu, Shiqian Zhao, Chenlong Yin, Jinhu Fu, Yibo Yan, Hanjun Luo, et al. 2025. A comprehensive survey in llm (-agent) full stack safety: Data, training and deployment. *arXiv preprint arXiv:2504.15585* (2025).
- [51] Chunqiu Steven Xia, Yinlin Deng, Soren Dunn, and Lingming Zhang. 2024. Agentless: Demystifying LLM-based Software Engineering Agents. *CoRR* (Jan. 2024). <https://openreview.net/forum?id=WYK8E00dm2>
- [52] Jinwei Xu, He Zhang, Yanjing Yang, Lanxin Yang, Zeru Cheng, Jun Lyu, Bohan Liu, Xin Zhou, Alberto Bacchelli, Yin Kia Chiam, et al. 2024. One Size Does Not Fit All: Investigating Efficacy of Perplexity in Detecting LLM-Generated Code. *ACM Transactions on Software Engineering and Methodology* (2024).
- [53] Zhou Yang, Zhipeng Zhao, Chenyu Wang, Jieke Shi, Dongsun Kim, Donggyun Han, and David Lo. 2024. Unveiling memorization in code models. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. 1–13.
- [54] Shenglai Zeng, Yaxin Li, Jie Ren, Yiding Liu, Han Xu, Pengfei He, Yue Xing, Shuaiqiang Wang, Jiliang Tang, and Dawei Yin. 2023. Exploring memorization in fine-tuned language models. *arXiv preprint arXiv:2310.06714* (2023).
- [55] Li Zhong and Zilong Wang. 2024. Can llm replace stack overflow? a study on robustness and reliability of large language model code generation. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 38. 21841–21849.
- [56] Albert Ziegler, Eirini Kalliamvakou, X Alice Li, Andrew Rice, Devon Rifkin, Shawn Simister, Ganesh Sittampalam, and Edward Aftandilian. 2024. Measuring github copilot's impact on productivity. *Commun. ACM* 67, 3 (2024), 54–63.