

**Document Version**

Final published version

**Licence**

Dutch Copyright Act (Article 25fa)

**Citation (APA)**

Koohestani, R., Bekker, P. D., Koç, B., & Izadi, M. (2025). Benchmarking AI Models in Software Engineering: A Review, Search Tool, and Unified Approach for Elevating Benchmark Quality. *IEEE Transactions on Software Engineering*, 52(2), 651-674. <https://doi.org/10.1109/TSE.2025.3644183>

**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.

# Benchmarking AI Models in Software Engineering: A Review, Search Tool, and Unified Approach for Elevating Benchmark Quality

Roham Koohestani , Philippe de Bekker , Begüm Koç , and Maliheh Izadi 

**Abstract**—Benchmarks are essential for unified evaluation and reproducibility. The rapid rise of Artificial Intelligence for Software Engineering (AI4SE) has produced numerous benchmarks for tasks such as code generation and bug repair. However, this proliferation has led to major challenges: (1) fragmented knowledge across tasks, (2) difficulty in selecting contextually relevant benchmarks, (3) lack of standardization in benchmark creation, and (4) flaws that limit utility. Addressing these requires a dual approach: systematically mapping existing benchmarks for informed selection and defining unified guidelines for robust, adaptable benchmark development. We conduct a review of 247 studies, identifying 273 AI4SE benchmarks since 2014. We categorize them, analyze limitations, and expose gaps in current practices. Building on these insights, we introduce BenchScout, an extensible semantic search tool for locating suitable benchmarks. BenchScout employs automated clustering with contextual embeddings of benchmark-related studies, followed by dimensionality reduction. In a user study with 22 participants, BenchScout achieved usability, effectiveness, and intuitiveness scores of 4.5, 4.0, and 4.1 out of 5. To improve benchmarking standards, we propose BenchFrame, a unified approach to improve benchmark quality. Applying BenchFrame to HumanEval yielded HumanEvalNext, which features corrected errors, improved language conversion, higher test coverage, and greater difficulty. Evaluating 10 state-of-the-art code models on HumanEval, HumanEvalPlus, and HumanEvalNext revealed average pass-at-1 drops of 31.22% and 19.94%, respectively, underscoring the need for continuous benchmark refinement. We further examine BenchFrame’s scalability through an agentic pipeline and confirm its generalizability on the MBPP dataset. Lastly, we publicly release the material of our review, user study, and the enhanced benchmark.<sup>1</sup>

## I. INTRODUCTION

**B**ENCHMARKS are essential for assessing artificial intelligence-driven software engineering (AI4SE) techniques. They provide standardized performance metrics, facilitate reproducibility, and guide innovation. However,

Received 7 March 2025; revised 26 November 2025; accepted 4 December 2025. Date of publication 16 December 2025; date of current version 13 February 2026. Recommended for acceptance by M. Zhou. (Corresponding author: Roham Koohestani.)

The authors are with the Electrical Engineering, Mathematics, and Computer Science (EEMCS), Delft University of Technology, 2628 XE Delft, The Netherlands (e-mail: r.koohestani@tudelft.nl).

Digital Object Identifier 10.1109/TSE.2025.3644183

<sup>1</sup><https://github.com/AISE-TUdelft/AI4SE-benchmarks>

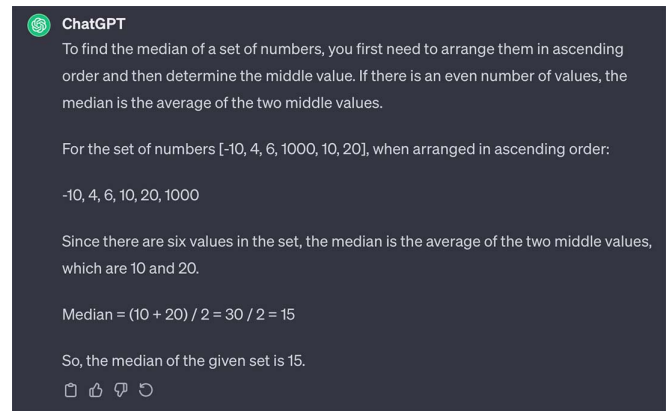


Fig. 1. ChatGPT repeats the same HumanEval error (captured on Dec 2023).

the exponential growth in benchmark development has introduced significant challenges: researchers must navigate an increasingly fragmented landscape to identify benchmarks that align well with their specific objectives. This complexity often incentivizes reliance on popular or widely-adopted benchmarks without scrutinizing their applicability, inherent limitations, or potential flaws. Such practices risk propagating biases, overestimating technical progress, and misdirecting research priorities.

A notable example of benchmark limitations in code generation evaluation is HumanEval [1], a widely-used dataset for assessing large language models (LLMs) such as Codex [1], Gemini [2], and GPT-4 [3]. HumanEval was downloaded 82 thousand times in July 2025 on a single platform (Hugging Face)<sup>2</sup>, reflecting the strong and sustained interest in this benchmark. Despite its broad adoption, HumanEval contains numerous flaws and inconsistencies [4]. For instance, Task 47, which requires computing the median of a numerical list, incorrectly states the median of the list  $[-10, 4, 6, 1000, 10, 20]$  is 15. When queried on this task, ChatGPT-3.5 Turbo reproduced the incorrect result (see Fig. 1). This suggests potential data contamination and benchmark overfitting which can artificially inflate performance scores.

As HumanEval remains a widely-used benchmark in both AI and software engineering communities, several efforts have

<sup>2</sup>[https://huggingface.co/datasets/openai/openai\\_humaneval](https://huggingface.co/datasets/openai/openai_humaneval)

sought to expand its language support [5], [6], [7], [8] or improve test coverage [9], [10]. However, these extensions often build upon the original dataset without addressing its fundamental deficiencies, allowing inherent issues, such as flawed canonical solutions, vague problem definitions, incorrect tests, and insufficient coverage to persist. Moreover, LLM-augmented improvements, such as automatic translation, often lack rigorous quality control as well. Lastly, as models have advanced, HumanEval and similar popular benchmarks have become increasingly saturated, with close to 100% scores for recent models.<sup>3</sup> This necessitates continuous elevation of the complexity of programs to better reflect the capabilities of contemporary models.

Our central objective is to **address the systemic challenges within AI4SE benchmarking**. We aim to address four key challenges in benchmarking: (1) the fragmentation of benchmark knowledge across tasks, (2) the difficulty of selecting contextually relevant benchmarks, (3) the lack of standardized approaches for benchmark creation and refinement, and (4) existing inherent flaws that limit benchmark utility. To address this, we use a three-stage approach where each stage builds upon the findings of the previous one.

More specifically, we first conducted a systematic review of benchmarking efforts in AI4SE from 2014 onward and identified **273** benchmarks from 247 studies. We then extracted a given benchmark's key metadata, including (1) objectives, (2) category, (3) programming language, (4) natural language, (5) relevant tasks, (6) extent to which tests are present, (7) scale of the dataset, (8) dataset source, (9) language specificity, (10) maintenance adequacy, (11) whether the dataset is reviewed or not, (12) whether it is frequently used or not, (13) the licensing, and (14) how it was created, to structure the AI4SE benchmark landscape. The findings from our review confirm a fragmented and difficult-to-navigate ecosystem, which motivated us to use these data to develop tooling to easily navigate the landscape of AI4SE benchmarking. **BenchScout**<sup>4</sup> is an *extensible semantic* search tool that enables users to efficiently identify relevant benchmarks for specific software engineering tasks. To build BenchScout, we applied clustering techniques to contextual embeddings derived from related studies and benchmark documentation, along with our manually-extracted metadata. Additionally, we employed dimensionality reduction techniques to visualize the AI4SE benchmark landscape. We conducted a user study with 22 participants from both industry and academia to gauge the *usability*, *effectiveness*, and *intuitiveness* of BenchScout. It achieved average scores of 4.5, 4.0, and 4.1 out of 5, respectively.

While BenchScout aids navigation, our analysis confirms that many foundational benchmarks are flawed, which means a systematic solution is required to deal with them. Based on the identified gaps and limitations in current benchmarks, we introduce **BenchFrame**, a peer-review-oriented methodology to improve the quality of the benchmarks of both existing and new benchmarks. To demonstrate its efficacy, we conduct a

case study on HumanEval and present *HumanEvalNext* as an enhanced version.

When evaluating performance using HumanEval (original) and HumanEvalNext (our improved version based on the BenchFrame), we observe a substantial decline in pass@1 accuracy. Across ten state-of-the-art open-source code models, the average pass@1 score decreases by 31.2%, with a median drop of 26.0%. Performance remains significantly lower even on HumanEvalPlus [10], an enhanced version of HumanEval, with an average decline of 19.94% in pass@1 scores. These results highlight the importance of continuously refining benchmarks to better guide future research and provide more realistic assessments of model performance.

In summary, our contributions are as follows.

- We conducted a comprehensive review of 247 studies, identified 273 AI4SE benchmarks, and analyzed their limitations and gaps (section III),
- We developed and released **BenchScout**, an extensible semantic search tool to facilitate locating appropriate AI4SE benchmarks. Our user study with 22 participants demonstrated its effectiveness (section IV),
- We propose a unified approach, **BenchFrame**, to improve the quality and reliability of benchmarks. A case study on HumanEval resulted in a refined, peer-reviewed benchmark, HumanEvalNext (section V). We assessed ten recent LLMs on this benchmark and verified significant drops in performance compared to the original HumanEval [1] and the enhanced version, HumanEvalPlus [10],
- We publicly share our data and details of the user study in our replication package.<sup>5</sup>

## II. RESEARCH QUESTIONS

To systematically address the challenges outlined in the introduction, we structure our study around our central research objective: **to diagnose the fragmentation and quality issues within the AI4SE benchmark ecosystem, provide practical tools to navigate this complexity, and establish a rigorous methodology for future quality improvement**. We deconstruct this objective into three guiding Research Questions (RQs), which are designed to be addressed sequentially.

- **RQ1: Diagnosing the problem:** To what extent is the AI4SE benchmarking landscape fragmented, and what are its primary, systemic limitations?
- **RQ2: Providing a Navigational Aid:** How can a semantic search and visualization tool improve the discoverability and selection of relevant AI4SE benchmarks?
- **RQ3: Proposing a Systemic Solution:** What constitutes a rigorous and scalable methodology for improving benchmark quality, and what is its measurable impact on model evaluation?

These research questions directly map to the three-stage approach of our study. **RQ1** fulfills the diagnostic part of our objective by mapping the landscape and identifying its core

<sup>3</sup><https://evalplus.github.io/leaderboard.html>

<sup>4</sup>Accessible through <https://evalpro.online/>

<sup>5</sup><https://github.com/AISE-TUDELfT/AI4SE-benchmarks>

problems. The findings from RQ1 motivate **RQ2**, which provides a practical tool to help the community navigate the benchmark fragmentation we uncovered. Finally, **RQ3** addresses the root cause of the quality issues found in RQ1 by proposing and validating a systematic approach for improvement, thereby fulfilling the final part of our research objective.

### III. EXISTING BENCHMARKS, A REVIEW

#### A. Search Criteria and Quality Assessment

We employed a systematic method to search for, identify, and classify AI4SE benchmarks. This procedure involved three primary stages: conducting structured searches, verifying credibility, and developing a taxonomy. We searched on two platforms, namely Google Scholar and Semantic Scholar, using the following keywords: “Benchmarks”, “Software Engineering”, “Large Language Models”, “Evaluation”, and “AI4SE”. We additionally searched for benchmarks present in the PapersWithCode datasets collection due to the popularity and wide usage of the platform.<sup>6</sup>

We selected these keywords based on a preliminary assessment of highly-cited benchmark research. Three authors worked together to revise the search criteria to ensure they were both relevant and complete. Our selection criteria targeted primary studies published in English from 2014 to 2025. After the initial paper collection process, duplicates were removed. Two authors reviewed the relevance of the identified pieces of literature, during which the originality (that is, the status as primary study), reproducibility, and accessibility of each study were accessed.

More specifically, for originality, we assessed whether a study was a primary study that introduced a new benchmark or a significant extension of an existing one. This filter was set in place to distinguish between benchmark-proposing papers from studies that merely applied them with a certain level pre-processing. In terms of Reproducibility and Accessibility, we evaluated whether the benchmark and its associated resources were made publicly available. We looked for the presence of a public repository (e.g., GitHub or HuggingFace) that would allow the community to access and use the benchmark. After this phase, we performed forward and backward snowballing. A total of 247 papers were found through this process.

After the papers were collected, the authors worked together to develop and consistently improve a taxonomy to efficiently categorize the benchmarks and extract metadata. Through continuous discussions, we identified initial essential categories while systematically gathering additional details for each study, including DOI and publication date. An iterative strategy was used to design the taxonomy and categorize SE tasks, starting with overarching categories such as reasoning, synthesis, and debugging. When the categories became too broad or lacked cohesion, the authors refined them into more detailed subcategories, establishing a multi-tiered hierarchy. Furthermore, it is important to highlight that during the categorization process, the

<sup>6</sup>This platform has since been discontinued and is no longer accessible. The previously used URL was <https://paperswithcode.com/datasets>




Fig. 2. Number of published benchmark papers (Jan 2014-Aug 2025).

two authors involved discovered 10 instances of disagreement over category classifications, resulting in an agreement rate of 96.4%. This outcome can be attributed to the relatively objective nature of benchmark categorization. Many of the papers either clearly specify their intended task or imply it within the text, with the latter being the source of discrepancies in the few instances where the raters disagreed.

In our evaluation, we not only considered benchmark datasets from academic sources, but also included those suggested by industry and subsequently adopted by scholars (for example, IBM CODAIT and Aider [11]). Our replication package also serves as a dynamic repository, allowing researchers to contribute additional benchmarks and related papers by submitting a pull request that includes the paper’s DOI.

After finalizing the taxonomy, and the list of benchmarks, we extracted (1) objectives, (2) categories, (3) programming language, (4) natural languages, (5) relevant tasks, (6) extent to which tests are present, (7) scale of the dataset, (8) dataset source, (9) language specificity, (10) maintenance adequacy, (11) whether the dataset is reviewed or not, (12) whether it is frequently used or not, (13) the licensing, and (14) how it was created, for each of the benchmarks.

#### B. Results of the Review

Our review revealed a significant increase in the number of benchmarks published over time (see Fig. 2). We identified 273 benchmarks in total, with 71 published in 2024 alone. Using an exponential projection, we estimate that this number will reach 109 for the year 2025. This highlights the growing impact of AI4SE benchmarking and the need for a comprehensive overview of the existing literature.

Additionally, upon analyzing the distribution of the tasks included in the review, we find that a great portion of benchmarks are related to *Code Generation*, more specifically, 34.4% of all the benchmarks included. Along with Code Understanding and Repair & Maintenance (with respectively 17.7% and 15.4%), these make up more than 67% of all benchmarks included. This, in itself, underscores a lack of attention to tasks such as vulnerability detection and code retrieval/search (cf. Fig. 3).

The data reveal a notable change over time in the distribution of research categories throughout the observed period (cf. Fig. 4). A key insight is the shift from code repair and maintenance to automated synthesis, marked by two opposing trends. Initially dominant, the Repair & Maintenance category

TABLE I  
OVERVIEW OF AI4SE BENCHMARKS STEMMING FROM HUMANEval[1] (PL DENOTES PROGRAMMING LANGUAGE, NL DENOTES NATURAL LANGUAGE)

Category	Name	Language(s)	# Tests
Original	HumanEval [1]	Python	Avg. 7.7
	MultiPL-HumanEval [6]	18 PLs	Avg. 7.7
	HumanEval-Fix [7]	6 PLs	Avg. 7.7
Improved Language Support	HumanEval-Explain [7]	6 PLs	Avg. 7.7
	HumanEval-Synthesize [7]	6 PLs	Avg. 7.7
	HumanEval-X [8]	5 PLs	Avg. 7.7
	Multi-HumanEval [5]	12 PLs	Avg. 7.7
	HumanEvalXL [12]	12 PLs, 23 NLs	Avg. 8.33
Improved Testing	mHumanEval [13]	Python, 204 NLs	Avg. 7.7
	HumanEval+ [10]	Python	Scaled $\times 80$
	HumanEval-MINI [10]	Python	Scaled $\times 47$
	HE-Eval [9]	Python	Scaled $\times 14$
Instruction-Based	InstructHumanEval <sup>7</sup>	Python	Avg. 7.7
Extended	EvoEval [14]	Python	Multiple categories, scaled with EVALPLUS

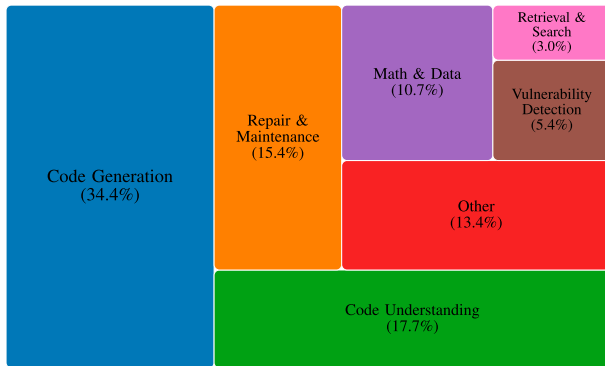


Fig. 3. High-level category distribution of included benchmarks (updated, alternating orientation).

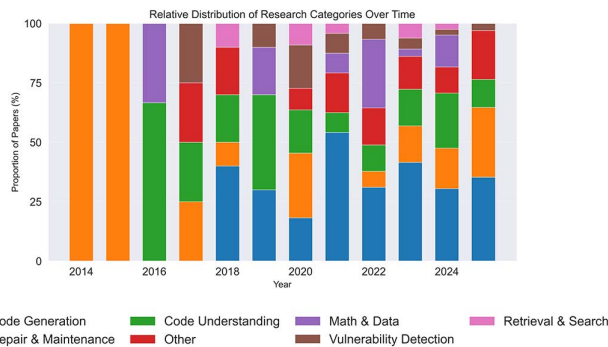


Fig. 4. Task distribution progression for the benchmarks included in the study.

gradually experiences a notable decline in its share. In contrast, starting around 2020, the Code Generation category gains prominence, which coincides with the advancement of large-scale language models. In the recent years presented, there is an emerging trend of thematic diversification; as the focus on Code Generation recedes from its high, categories like Code Understanding and Retrieval & Search show a relative resurgence. This suggests a community-wide transition towards tackling

TABLE II  
OVERVIEW OF AI4SE BENCHMARKS DERIVED FROM MBPP[15]

Category	Name	Language(s)	# Problems
Original	MBPP [15]	Python	974
Improved Language Support	MultiPL-MBPP [6]	18 PLs	354-397 per PL
	MBXP [5]	13 PLs	848-974 per PL
Improved Testing	MBPP+ [10]	Python	427
	MBPP-Eval [9]	Python	974

additional challenges related to verification, integration, and the understanding of algorithmically generated outputs. In the following, we summarize key insights from our review.

**HumanEval Benchmark Family:** Currently, one of the most popular AI4SE benchmarks is HumanEval [1], used to evaluate the performance of many notable code-aware models (e.g., Codex [1], Gemini [2], and GPT-4 [3]). This benchmark is used mainly for code synthesis, though there exist some variations for code repair and code explanation [7]. Table I presents the family of HumanEval benchmarks. After an in-depth analysis of these benchmarks, we identified the following issues: (1) incorrect tests, (2) lack of proper test coverage, (3) incorrect canonical solutions, and (4) imprecise problem definitions. While there are versions that have improved the language support [5], [6], [7], [8] and test coverage [9], [10], there is no version that contains all the improvements combined nor fixed the original issues. The issues for enhancing the original dataset can be generalized as follows:

- Variants that cover multiple languages have duplicated the original issues.
- Variants that added tests used the original incorrect solutions to generate the output.
- Variants based on human corrections or translations are inconsistent.

Furthermore, production systems like ChatGPT-3.5 tend to replicate errors from the original HumanEval benchmark. This indicates potential contamination from the benchmark data,

<sup>7</sup><https://huggingface.co/datasets/codeparrot/instructhumaneval>

TABLE III  
OVERVIEW OF COMPETITIVE PROGRAMMING, CODE COMPLEXITY, AND CODE EFFICIENCY BENCHMARKS

Category	Name	Language(s)	# Tests
Competitive Programming	CodeContests [16]	12 PLs	Avg. 203.7
	APPS [17]	Python	Avg. 13.2
	LiveCodeBench [18]	Python	Avg. 17.23
	LeetCode [19]	Python	Avg. 135
	CodeElo [20]	N/A	408 problems
	FVAPPS [21]	Python	4715 problems
	KodCode [22]	12 PLs	Avg. 7.52
	FC2Code [23]	Python	320 flowcharts
	CodeContests+ [24]	Python	11,690 problems
Code Complexity	LiveCodeBenchPro [25]	Python	584 problems
	CoRCoD [26]	Java	932
	GeeksForGeeks (GFG) [27]	C++, Python	$\pm 1,400$ per lang.&categ
	CODAIT <sup>8</sup>	Python	4,000,000
	CodeComplex [28]	Java, Python	4,900 per language
Code Efficiency	PythonSaga [29]	Python	185
	EffiBench [30]	Python	Self-defined, avg. 100
	CODAL [31]	Python	3 ref. / problem
	PIE [32]	C++	82.5(median, train)
	COFFE [33]	Python	756

TABLE IV  
OVERVIEW OF DATA SCIENCE & DOMAIN-SPECIFIC BENCHMARKS. (JN REFERS TO JUPYTER NOTEBOOKS.)

Name	Language(s)	# Tests	Comment
DS-1000 [34]	Python	Avg. 1.6	7 DS/ML libraries
NumpyEval [35]	Python	Avg. 20 functions	NumPy (101 problems) (Avg. 1 variable)
PandasEval [35]	Python	Avg. 20 functions	Pandas (101 problems) (Avg. 1 variable)
JuICe [36]	Python, JN	N/A	Cell completion (1.5M/3.7K train test)
DSP [37]	Python, JN	Available	Cell completion (1,119 problems)
ExeDS [38]	Python, JN	Execution Based	Cell generation (ground truth), 534 tasks
DSEval [39]	Python	custom approach	Models Evaluated via the DSEval Approach from the Paper
TorchDataEval [40]	Python	50	Private PyTorch data library
MonkeyEval [40]	Python	101	Private Pandas library fork
BeatNumEval [40]	Python	101	Private NumPy variant
Bio-Coder [41]	Python	1,026	Identify and import necessary classes for given task
Bio-Coder-Rosalind [41]	Java	1,243	
	Python	253 golden solution	Generate code for question evaluates whether a model can generate React web-app
WebApp1k [42]	React	Available	

with the systems not only producing incorrect answers but also seemingly optimizing to match the flawed responses from the widely used benchmark. Given the widespread use and ongoing popularity of the HumanEval benchmark in the research community, it is crucial to address its inherent flaws to prevent the perpetuation of these issues.

**MBPP Benchmark Family:** Another AI4SE benchmark, highly similar in style and popularity compared to HumanEval, is MBPP [15]: Mostly Basic Python Problems. It contains nearly a thousand crowdsourced problems, where almost half of it is sanitized and separately released. Furthermore, several enhancements have been published for MBPP (Table II). Upon a more in-depth analysis of MBPP and its family of benchmarks,

there are many signs suggesting deficient quality. One notable problem is the lack of proper testing, as MBPP originally only has three (rather trivial) tests per problem, which are all revealed in the prompt as well. With such a test suite in place, evaluation metrics become unstable and insignificant for proper comparison. The strength of the written tests and solutions themselves is not only troublesome in the original data but also the sanitized data features many flaws (even in *corrected* variants [10]). From negligible observations such as poor syntax (e.g., too many spaces, Python method names starting with a capital – this is a common convention to only use for classes) to uncaught bugs and edge cases that break the implementation. While there are enhancements that improve the language support and extend the test cases, they are all built upon inadequate foundations, which renders any MBPP benchmark suboptimal for proper assessment.

<sup>8</sup>CODAIT-2021 <https://ibm.co/4emPBla>

TABLE V  
OVERVIEW OF MATHEMATICAL REASONING BENCHMARKS

Name	Language(s)	# Problems
MATH [43]	English	12,500
MATH500 [44]	English	500
MathQA [45]	English	37,297
MathQA-Python [15]	Python	23,914
MathQA-X [5]	Python, Java, JS	1,883 per language
LiLA [46]	Python	133,815 questions 358,769 programs
MultiArith [47]	English	600
GSM8K [48]	English	1,320
GSM-HARD [48]	English	1,320
TheoremQA [49]	English	800
PECC [50]	Python	1,006
BRIGHT [51]	English	395
AMC12 <sup>9</sup>	English	82

TABLE VI  
OVERVIEW OF NATURAL LANGUAGE BENCHMARKS. THESE ARE SPECIFICALLY THE TEXT2CODE BENCHMARKS

Name	Language(s)	# Problems
CoNaLa [52]	English → Python	2,879
MCoNaLa [53]	Spanish, Japanese, Russian → Python	896
CoNaLa-SO [54]	English → Python	10,000
APPS [17]	English → Python	10,000
APPS-Eval [9]	English → Python	10,000
AixBench [55]	English, Chinese → Java	175
Natural2Code [2]	English → Python	Unknown
CoSQA [56]	English → Python	20,604
WebQueryTest [57]	English → Python	1,046
AdvTest [57]	English → Python	280,634
CONCODE [58]	English → Java	104,000
MTPB [59]	English → Python	115
CAASD [60]	English → Python	72
Shellcode_IA32 [61]	English → IA32/Shell	3200
Odex [61]	Spanish, Japanese, Russian, English → Python	945 90, 164, 252, 439 1707 test total
PSB2 [62]	English → Clojure, Python	25 question-answer pairs
TACO [63]	English → Python	1,539,152 on 26,433 distinct tasks
Turbulence [64]	English → Python	60 (with 420 total test cases)
Aider <sup>10</sup>	English → C++, GO, Java, JS, Python, Rust	225 problems
NL2ML-lib [65]	English → Python	11,000
RMCBench [66]	English → 9 Languages	473 malicious prompts
Evil [67]	English → Python, IA_32	19255
Exec-CSN [68]	English → Python	1,931
CodeIF [69]	English → Java, Python, Go, C++	1,200 tasks
CodeIF-Bench [70]	English → Python	122 tasks (42 repositories)
ARCADE [71]	English → Python	1082
StackEval [72]	English → Python	925
SwiftEval [73]	English → Swift	28
CoSQA+ [74]	English → Python	412,080 pairs

**Other Existing Benchmarks:** Besides HumanEval and MBPP, the standardized benchmarks for code synthesis evaluation, there are many more considerable benchmarks for

<sup>9</sup><https://huggingface.co/datasets/AI-MO/aimo-validation-amc>

<sup>10</sup><https://github.com/Aider-AI/aider/blob/main/benchmark/README.md>

TABLE VII  
OVERVIEW OF NATURAL LANGUAGE BENCHMARKS (CONTINUED). THE FIRST TWO BENCHMARKS ARE TEXT2TEXT (RELATED TO CODE) BENCHMARKS, WHILE THE REST ARE CODE2TEXT BENCHMARKS

Name	Language(s)	# Problems
InfCoder-Eval [75]	English → English	270
BRIGHT [51]	English → English	1,398
DeepCom [76]	Java → English	588K
Hybrid-DeepCom [77]	Java → English	466k
BinSum [78]	Binary functions → English	557K
Code Attention [79]	Java → English	11 projects
Funcom [80]	Java → English	2.1M problems
CodeSum [81]	Java → English	410,630
CoDesc [82]	Java → English	4.21M datapoints
Parallel [83]	Python → English	150k function/doc pairs
CoDocBench [84]	Python → English	4573 code/doc pairs
PoorCodeSum [85]	Java, Python, Go → English	10,955, 14,918, 8,122
P-CodeSum [86]	Python, Java, Go, JS, PHP, Ruby → English	1,500 pairs

TABLE VIII  
OVERVIEW OF SQL-RELATED BENCHMARKS

Name	Language(s)	# Problems
BIRD [87]	English → SQL	12,751
KaggleDBQA [88]	English → SQL	272, paired with golden solutions
StacQc [89]	English → Python/SQL	147,546 / 119,519 question-answer pairs
Spider(V2 <sup>11</sup> ) [90]	English → SQL	632 queries
Spider-Syn [91]	English → SQL	(7000 / 1034)
Spider-Real [92]	English → SQL	508
Spider-DK [93]	English → SQL	535 pairs
Spider-CN [94]	Chinese → SQL	9691 queries
SParC [95]	English → SQL	4,298 question sequences
Lyra [96]	English, Chinese → python, SQL	2000
DuSQL [97]	Chinese → SQL	23,797 question/SQL pairs
CoSQL [98]	English → SQL	3,007 Question Sequences
SynSQL-2.5M [99]	English → SQL	2,544,390
PAUC [100]	Russian → SQL	9,691
Ar-Spider [101]	Arabic → SQL	9,691
Tur2SQL [102]	Turkish → SQL	10,809 question/SQL pairs

assessing various categories of SE tasks. We share several additional categorized tables as a guide for finding specific AI4SE benchmarks and highlight the most notable benchmarks for each in detail.

Table III Features benchmarks with competitive programming as their root, i.e., those used for understanding code complexity and efficiency. Table IV features a set of benchmarks specifically designed to evaluate the performance of models on data science-related tasks along with some other domain-specific SE tasks. A notable example in this table is the bio-coder series of benchmarks specifically designed for bioinformatics tasks. To assess the mathematical reasoning capabilities of ai4se models, see Table V. Besides numbers and code, natural language is also a key component in AI4SE. From supporting instruction-tuned AI4SE models, which align more with the human brain [103], that aim to accomplish question and answering

<sup>11</sup>see Table XIII

TABLE IX  
OVERVIEW OF SELECTED REAL-TO-LIFE SE BENCHMARKS. (NOTE: X/Y/Z DENOTES TRAIN/DEV/TEST)

Category	Benchmark	Language(s)	# Problems
Software Development & Agent Benchmarks	DevBench [104]	Python, C/C++, Java, JS	22 repositories
	DevEval [105]	Python	1,874
	CoderUJB [106]	Java	2,239
	CODAL [31]	Python	500
	ToolQA [107]	Python, Math, English	800(Easy)/730(Hard)
	MIT [108]	Python, English	586 Problems
	SAFIM [109]	Python, Java, C++, C#	17720
	AgentBench [110]	N/A	1360 prompts
	CSR-Bench [111]	Python	100 repositories
	Class Level	ClassEval [112]	Python
CONCODE [58]		English, Java	104,000
BigCodeBench [113]		Python	1,140
OOP-Bench [114]		Python	431
CodeSense [115]		Python, C, Java	2125 (Python), 876 (C) 875 (Java)
ClassEval-T [116]		Java, C++	1,243
SWE-bench [117]		Python	19,008 (Train), 225 (Dev), 2,294 (Test) 144 (Small)
CrossCodeEval [118]		C#, TS, Java, Python	2,665 (Python), 2,139 (Java), 3,356 (TS), 1,768 (C#)
CoderEval [119]		Java, Python	230
DotPrompts [120]		Java	105538 problems (1420 methods)
Project & Cross-File	BigCloneBench [121]	Java	25,000 Java Systms
	DI-Bench [122]	Python, C#, Rust, JS	581 repositories (w/ dependencies)
	DyPyBench [123]	Python	50 repositories
	Multi-SWE -bench [124]	Java, TS, JS Go, Rust, C, C++	500 (Python), 128 (Java), 224 (TS), 356 (JS), 428 (Go), 239 (Rust) 128 (C), 129 (C++)
	KernelBench [125]	Python	250
	CodeMEnv [126]	Python, Java	587 (Python), 335 (Java)
	CodeEditorBench [127]	Python, Java, C++	7,961
	ProjectEval [128]	Python	284

(QA) similar to the widely recognized platform StackOverflow, to summarizing code and generating tags, Table VI and Table VII feature AI4SE benchmarks including natural language: text-to-code, code-to-text and text-to-text (code related). We have intentionally isolated all SQL-related benchmarks due to their abundance and to facilitate locating the correct benchmark; presented in Table VIII.

While translating natural language is more trivial nowadays, translating code remains challenging due to various reasons (e.g. versioning, semantics, dependencies). With the lack of diversity in language support for AI4SE benchmarks and also benefiting numerous other SE tasks, Table XIV features an overview of resources that can support the ongoing development of code translation.

However, the workflow of a developer is not merely an exercise of writing snippets of code for given descriptions, but rather having a general overview of a project and how one can implement functionality such that it fits well into a collective code base. For this, Table IX and Table X provide a collection of benchmarks that examine the capabilities of models to generate code on a larger scale.

The utilization of APIs plays a significant role in AI4SE benchmarks, specifically for models with Retrieval Augmented Generation (RAG) capabilities. In Table XI, prominent benchmarks focusing on leveraging the power of APIs are

denoted. Furthermore, Table XII lists benchmarks related to pseudocode, followed by an overview of notable crowd-sourced AI4SE resources in Table XIII.

With AI4SE models mainly being utilized for program synthesis, it remains relatively questionable how effective these models are in generating tests and repairing bugs, as it is unclear whether these models *truly* understand code. For example, Siddiq et al. [232] observed Codex [1] being able to get above 80% coverage for HumanEval [1], yet many test smells were discovered and for another dataset, no higher than 2% coverage was attained. This reveals the importance of benchmarking AI4SE models' capabilities in test generation, bug repair, and understanding. In Table XV, several benchmarks are listed that make an effort to assess the aforementioned. Additionally, Table XVI table features benchmarks that have been designed to evaluate a model's capabilities in dealing with everyday tasks of a software engineer (e.g., merge-conflict repair, Code Reviews, etc.).

Additionally, we collect a set of generic benchmarks in Table XVII. These benchmarks are not only single-task benchmarks like others previously seen in this section but are rather a collection of tasks spanning a wide range of languages and task types. A notable, and widely known benchmark in this category is the BigBench Benchmark [233] which consists of 167 tasks (not all relevant to AI4SE).

TABLE X

OVERVIEW OF SELECTED REAL-TO-LIFE SE BENCHMARKS. THIS TABLE CONTAINS THE REPOSITORY-LEVEL BENCHMARKS. (CONTINUED - NOTE: X/Y/Z DENOTES TRAIN/DEV/TEST)

Benchmark	Language(s)	# Problems
RepoBench [129]	Python, Java	Cross-file: 8,033
		In-file: 7,910
RepoEval [130]	Python	1,600 (line), 1,600 (API) 373 (function)
EvoCodeBench [131]	Python	275
SketchEval [132]	Python	19 repositories (5 easy, 8 medium 6 hard)
Stack Repo [132]	Python	(435,890 / 220,615 / 159,822) answer pairs
ML-BENCH [133]	Python & Bash	9641 problems
CodeGen4Libs [134]	Java	403,780 prompts
SWE-rebench [135]	Python	294 tasks, 169 repositories
	Java	165 (Java)
SWE-Polybench [136]	JS	1017 (JS),
	TS	729 (TS)
	Python	729 (199 (Python))
HumanEvo [137]	Python, Java	200 (Python), 200 (Java)
REPOCOD [138]	Python	980 functions (11 projects)
FEA-Bench [139]	Python	83 repositories
JavaBench [140]	Java	389 methods (106 classes)
SWE-bench live [141]	Java	1,319 tasks (93 repositories)
SWE-Lancer [142]	N/A	1,488 tasks

Finally, we have identified a rather large count of papers that have to do with the matter of Log Parsing and Log Statement generation. we have presented these benchmarks in Table XVIII.

### C. Limitations of Existing Benchmarks

Across our review, we observed several trends in the existing AI4SE benchmarks. The majority are language-specific, most commonly targeting Python or Java. Although these languages are popular in both education and industry, this narrow scope potentially limits their generalization across domains and programming environments. Many benchmarks suffer from poor maintenance, with little to no updates or active support following their initial release. Official and dynamic leaderboards are often missing, which makes it difficult to fairly compare model performance across time or publications, particularly as newer models continue to emerge. Notably, benchmark popularity does not appear to strongly correlate with its maintenance quality or whether it was peer reviewed, suggesting other factors such as visibility or convenience may drive adoption.

There is also variation in how benchmarks are constructed. Manually created datasets are often significantly smaller, whereas newer benchmarks increasingly leverage LLMs to generate large volumes of tasks. Mined datasets which are mostly sourced from GitHub, can be large but are frequently under-specified, with limited transparency into the repositories used or the sampling heuristics applied.

Table XIX Summarizes these limitations and provides representative benchmarks for each. The full list of evaluated

benchmarks and their metadata can be found in the replication package of the study.

In conclusion, based on our in-depth inspection of HumanEval and MBPP and combined with the inspection of other benchmarks in our review, we obtained an overview of the limitations in current AI4SE benchmarks which we use as a guide in shaping our methodology proposed in section V.

## IV. BENCHSCOUT- LOCATING AI4SE BENCHMARKS

Due to the abundance of AI4SE benchmarks, identifying the most suitable one for a specific SE task can be challenging. As a result, many default to evaluating their models on popular benchmarks like HumanEval [1] which has its own flaws.

To address this gap, we developed BenchScout<sup>12</sup>, a tool to systematically and semantically search the existing benchmarks and their corresponding use cases. We additionally provide an interface to visually evaluate the closeness and similarity of a group of datasets, along with capabilities to find relations between citing bodies for identifying patterns relevant to different use cases.

### A. Context Extraction and Visualization

1) *Overview of the Semi-Automated Pipeline:* To effectively contextualize and visualize the growing corpus of AI4SE benchmarks, we developed a semi-automated pipeline that extracts, embeds, clusters, and visualizes benchmark metadata. As illustrated in Fig. 5, the system is organized into four modular layers: data collection, metadata enrichment, clustering and labeling, and front-end interaction.

The pipeline begins with benchmark sources collected from a curated spreadsheet, user-submitted suggestions, and an automated paper discovery script using the Semantic Scholar API to make suggestions. All inputs are aggregated into a processing queue, where metadata is enriched using both automated API calls and basic scraping techniques (e.g., fetching GitHub or Hugging Face README files). These metadata are transformed into dense vector representations using OpenAI's `text-embedding-3-small` model.

To support user navigation and cluster discovery, the high-dimensional embeddings are reduced via UMAP and grouped using HDBSCAN. A key component of the pipeline is the integration of GPT-based cluster labeling, which generates short, descriptive names for each cluster based on the titles and descriptions of the papers it contains. This step is fully automated but subject to human validation when new clusters emerge; meaning, these clusters *can* be modified by humans if needed.

2) *Balancing Automation and Oversight:* Unlike other prior approaches that rely on manual curation or entirely unsupervised grouping, our pipeline adopts a semi-automated architecture designed to scale with minimal human intervention while retaining curatorial oversight. Specifically, the system regularly ingests new suggestions, both from automated Semantic Scholar queries and from users interacting with the front-end

<sup>12</sup><https://evalpro.online/search.html>

TABLE XI  
OVERVIEW OF SELECTED API AND RETRIEVAL BENCHMARKS BY CATEGORY

Category	Benchmark	Sources/API(s)	# Problems
API Prediction	RestBench [143]	Spotify, TMDb	57, 100
	APIBench-Q [144]	StackOverflow, Tutorial Websites	6,563 (Java), 4,309 (Python)
	BIKER [145]	StackOverflow	33,000
	Gorilla APIBench [146]	HuggingFace, TensorHub, TorchHub	925, 696, 94
	Gorilla APIZoo [146]	Open submissions (Google, YouTube, Zoom, etc.)	–
Retrieval & Planning	API-Bank [147]	73 commonly used APIs	753
	CodeRAG-Bench [148]	Competition solutions, tutorials, documentation, StackOverflow, GitHub	25,859
	Search4Code [149]	Bing	6596(java)/4974(c#)
	CoIR [150]	GitHub, StackOverflow, and Various Benchmarks	2.38M (corpus) 3.37(queries)
Memorization	SATML-ext [151]	GitHub	1,000 samples
API Misuse Detection	ExampleCheck [152]	StackOverflow	100 (Java)
	ROBUSTAPI [153]	StackOverflow	1208 (18 Java APIs)
	APIMU4C [154]	Juliet Test Suite, ITC, OpenSSL, Curl, Httpd	2272 (C)

TABLE XII  
OVERVIEW OF AI4SE BENCHMARKS RELATED TO PSEUDOCODE

Category	Benchmark	Language(s)	# Problems	Crowdsourced
Pseudocode to Code	SPoC [155]	C++	18,356	Yes
	NAPS [156]	Java/UAST	17,477	No
	PseudoEval [157]	Python, C++, Rust	1,060	No
Code to Pseudocode	Django [158]	Python, English & Japanese	18,805 (Train), 1,000 (Dev), 1,805 (Test)	No

TABLE XIII  
OVERVIEW OF SELECTED CROWD-SOURCED BENCHMARKS (NL  
DENOTES NATURAL LANGUAGE)

Benchmark	Language(s)	# Problems	Source
WikiSQL [159]	NL → SQL query	80,654	Amazon MTurk (2017)
Spider [160]	NL → SQL query	10,181	11 Yale students (2018)
NL2Bash [161]	NL → Bash	9,305	Upwork (2018)
NAPS [156]	Java/UAST → Pseudocode	7,477	Self-hosted crowdsourcing, programming community (2018)
SPoC [155]	C++	18,356	programming websites (2019)
MBPP [15]	Python	974	Google Research, internal crowdworkers (2021)

interface, and proposes candidate entries for including in the pool of benchmarks. Although metadata extraction and embedding are fully automated, curators retain the ability to approve or reject new data sources, review GPT-generated labels, and fine-tune cluster boundaries as needed.

3) *Interactive Visualization*: The final layer of the pipeline provides an interactive 2D interface that enables users to explore benchmarks based on their semantic similarity. The visualization supports fuzzy search, local citation graphs, and interactive

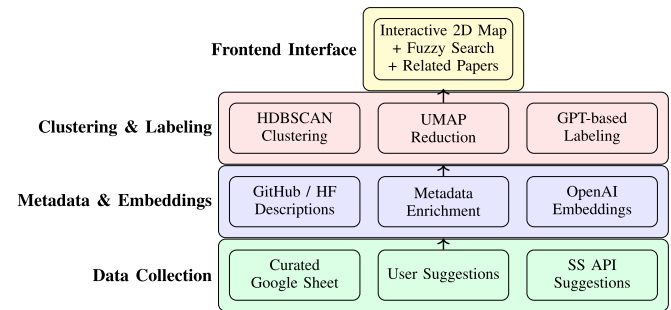


Fig. 5. Pipeline architecture of BenchScout. This includes automated metadata extraction, semantic embedding, clustering, and interactive visualization.

paper details. Importantly, user interactions on the frontend, such as suggesting new benchmarks, are routed back into the data collection layer (cf. Fig. 5).

### B. Additional Features

To improve the search and exploration experience, we incorporated several key functionalities. First, a **text-based search interface** allows users to find articles by title and abstract using **fuzzy search** via the Fuse.js library<sup>13</sup>. Users can also issue **advanced field-specific queries** (e.g., `language:python`) or **exact phrase queries** (e.g., “code generation”), and benefit from **autocomplete suggestions** for common terms, programming languages, tasks, and datasets.

<sup>13</sup><https://www.fusejs.io/>

TABLE XIV  
OVERVIEW OF PROGRAMMING LANGUAGE TRANSLATION BENCHMARKS (NOTE: X/Y/Z DENOTES TRAIN/DEV/TEST)

Category	Name	Language(s)	# Samples
Programming Languages	CodeTrans [57]	C#, Java	11,800
	TransCoder-ST [162]	C++, Java, Python	437,030
	CoST [163]	7 programming languages	16,738
	AVATAR [164]	Java, Python	7,133 / 476 / 1,906
	Multilingual-Trans [165]	8 programming languages	30,419 total
	NicheTrans [165]	Various niche languages	236,468 total
	LLMTrans [165]	8 programming languages	350
	G-TransEva [166]	5 programming languages	400 total
	CODEEDITOR [167]	C# & Java	6613
	RustRepoTrans [168]	C++, Java, Python → Rust	375
	AVATAR-TC [169]	Java, Python	55,179 / 443 / 1,746
	RustRepoTrans [168]	C++, Java, Python → Rust	375
RepoTransBench [170]	Python → Java	100 repositories	
Libraries	DLTrans [165]	PyTorch, TensorFlow, MXNet, Paddle	408 total
Intermediate Representation	SLTrans [171]	14 Languages → LLVM-IR	4M
Language Conversion Frameworks	MultiPL-E [6]	19 programming languages	-
	MultiEval [5]	13 programming languages	-

A dynamic **2-D UMAP visualization** displays papers embedded in semantic space. Hovering reveals abstracts via a **Paper Content Tooltip**, while double-clicking redirects to the paper’s DOI page. Paper nodes are sized based on **normalized citation count**, colored by cluster, and annotated with metadata such as year, tasks, and dataset names.

Clicking a point activates the **Related Papers feature**, which shows (1) The top-5 most similar papers based on cosine similarity, (2) A **detailed paper overview** with title, abstract, authors, venue/journal, and publication type, (3) A list of **citing papers** with sortable tables and filters by year and venue, and (4) An interactive **Paper Citations Graph**, which visualizes inter-citation relationships among citing papers.

Users can navigate between **chart**, **list**, and **grid** views, paginate results, and apply **advanced filters** by programming/natural language, dataset, task, cluster, year, and citation count. A **search history** and **keyboard shortcuts** improve usability. The entire interface supports responsive interaction with detailed feedback and error handling.

### C. User Study

With all the mentioned features, we aim to create a platform that can be extended and used by both academics and practitioners alike to find the appropriate dataset/benchmark for their use case with more ease. To evaluate how effective and usable this new tool is for the end-users, we conducted a user study on 22 people from both industry (9) and academia (13). In selecting demographics for the BenchScout user study, we aimed to assess the tool’s effectiveness across a diverse group of users with varying degrees of expertise. This approach seeks to determine the tool’s applicability for individuals at either end of the spectrum; whether they are beginners or seasoned experts, from academia or industry. The ultimate objective of the tool is to facilitate the selection of the appropriate benchmark, making it more accessible irrespective of prior knowledge. For this, we had each participant interact with the tool for however long they

saw fit and asked them to fill out a questionnaire consisting of eleven 5-point Likert Scale questions and three open questions.

1) **Questionnaire Design:** The questionnaire designed for the study was divided into four key sections to gather feedback about the tool, namely the participant’s background, the quality of the search functionality, the quality of the cross-referencing feature, and the overall user experience. Table 20 is an overview of the posed questions.

2) **Results and Analysis:** We analyzed the data collected through the questionnaire both quantitatively (Likert scale responses, scaling between 1-5) and qualitatively (open-ended questions). The detailed results are provided in the replication package.<sup>14</sup> Below, we present an overview.

The respondents of the questionnaire were generally familiar with AI4SE, with an average familiarity rating of 3.8. There were varying levels of experience in the field, with the range between 1-3 years being most common (eight people). More concretely, in our pool of participants, the demographic distribution was as follows:

- 1) **Roles:** The participants held diverse roles, including 6 Researchers, 5 PhD Candidates, 5 Students, 4 Software/Research Engineers, and 2 Lead Researchers
- 2) **Experience Level:** 8 participants had 1-3 years, 6 had 3-5 years, 3 had 5+ years, and 5 had less than 1 year of experience
- 3) **Familiarity with AI4SE:** On a scale of 1 (Not familiar) to 5 (Very familiar), 8 participants rated themselves a “5” and 7 rated themselves a “4,” indicating a strongly informed participant pool

In terms of search functionality, the tool scored high on usability with an average rating of 4.5, showing that users found it easy to navigate. The intuitiveness of the search interface received a solid 4.1, while its effectiveness in helping users find benchmarks was rated 4.0. The visual evaluation feature received 3.8, indicating some room for improvement in how visual elements assist in the search process.

<sup>14</sup><https://github.com/AISE-TUdelft/AI4SE-benchmarks>

TABLE XV  
OVERVIEW OF AUTOMATED PROGRAM REPAIR, FAULT LOCALIZATION, AND  
VULNERABILITY DETECTION BENCHMARKS.

	Benchmark	Language(s)	#Samples
Automated Program Repair & Fault Localization	Defects4J [172]	Java	835
	GitBug-Java [173]	Java	199
	EvalGPTFix [174]	Java	4530
	TutorCode [175]	C++	1239
	GHRB [176]	Java	107
	IntroClass [177]	C	998
	ManyBugs [177]	7 Languages	185
	DebugBench [178]	C++, Java Python	1,438 & 1,401 & 1,414
	QuixBugs [179]	Java	40 (locations of bugs)
	RES-Q [180]	Python, JS	100 hand-crafted questions + tests
	StudentEval [181]	Python	1,749 buggy programs (48 * 3 tests)
	Re-Factory [182]	Python	1783(buggy) 2442(correct)
	ConDefects [183]	Python	526(Python)
	Java Cerberus [184]	Java(477) C, C++, Java	2242 (across 4 tasks)
Vulnerability Detection	RepairBench [185]	Java	574
	MaRV [186]	Java	693
	BugsinPy [187]	Python	493 bugs
	CVEFixes [188]	Various	5,365
	LLMSEval [189]	C	150 (on 25 vulnerabilities)
	SecurityEval [190]	6 languages	130 (on 75 vulnerabilities)
	Vul4J [191]	Java	79 vulnerabilities
	FormAI [192]	C	112k instances
	VJBench [193]	Java	42 vulnerabilities
	SmartBugs [194]	Solidity	69 Vulnerable Smart Contracts
	Devign [195]	C	4 large Software Repos
	D2A [196]	C/C++	6 OSS Programs
	BigVul [197]	C/C++	348 Projects
	SARD <sup>15</sup>	Java, C, C++	32k <sup>16</sup>
	C#, PHP Juliet 1.3 <sup>17</sup>	C/C++	64k <sup>18</sup>
	NVD <sup>19</sup>	Various	265k <sup>20</sup>
	ARVO [198]	C, C++	1,001 vulnerabilities
	VADER [199]	15 languages	174 vulnerabilities
	ManyVuls4J [200]	Java	103 vulnerabilities
Software Testing	CoverageEval [201]	Python	1160
	ATLAS [202]	Java	9,275 projects
	HITS [203]	Java	10 projects
	MeMo [204]	Java	9 projects
	MLAPIs [205]	Python	63 applications
	CoderUJB [206]	Java	2,239
	TestBench [207]	Java	108
	TestEval [208]	Python	210
	TARBENCH [209]	Java	45,373 (59 projects)
	ProjectTest [210]	Python, Java, JS	20 per language
	CLOVER [211]	Python	845

When evaluating the cross-referencing features, respondents found the overall usefulness to be 4.1. However, the tool's ability to help users understand connections between different benchmarks was at 3.7. The visual interface for exploring these connections was rated 3.9, suggesting that while users

found it generally helpful, enhancements could improve its utility.

Regarding the overall user experience, participants gave an average rating of 4.2, with a score of 4.0 on the likelihood of using the tool in their own research. Several issues were highlighted, particularly around the dimensionality reduction and how the scatter plot is organized and presented. Users also noted that the citation network feature becomes less effective with larger papers and called for improved clustering by topic and additional features to explain and control visualizations.

The participants' responses confirmed our findings and highlighted the lack of a specific tool dedicated to locating AI4SE benchmarks. Instead, respondents commonly rely on generic platforms like Huggingface, Semantic Scholar, Google, and ConnectedPapers. When compared to these tools, BenchScout received an average score of 4.2 out of 5, with 5 indicating a much better experience. One participant mentioned using their personal network to find benchmarks, which limits broader access, further supporting the need for the proposed tool.

The tool's ability to meet the professional needs of users was rated 4.2 which affirmed its usefulness in the AI4SE domain. However, respondents suggested several additional features that could enhance its functionality, such as pagination for citations, incorporating metadata and additional information about the papers in the search process, improved clustering and filtering options, and sorting citations based on specific criteria. Additional requests included dark mode support, better overall search functionality, and clearer explanations and control over the chart visualizations. Based on the users' feedback and components' scores, we prioritized these features and incrementally added them to the platform.

In conclusion, the tool is largely perceived as useful and user-friendly, though several areas, particularly around visualization, citation handling, and filtering options, could be improved to enhance the overall user experience and the tool's effectiveness. Due to time constraints, we prioritized and implemented key features, leaving some for the future.

## V. BENCHFRAME

In this section, we propose BenchFrame, a detailed and peer-review-oriented approach for improving the quality of existing benchmarks. We explain our approach through a case study in which we propose HUMANEVALNEXT, a corrected foundation based on the HumanEval benchmark.

The HumanEval benchmark has long been the de facto standard for evaluating the code-generation capabilities of AI4SE models. It has recently been used to evaluate the latest and greatest LLMs from large companies such as Google Gemini [2] and OpenAI GPT4 [3]. Despite the great fame, however, our review in section III points towards the existence of numerous notable problems with this foundational benchmark, namely,

<sup>15</sup><https://samate.nist.gov/SARD/>

<sup>16</sup>As of 4th Feb 2025

<sup>17</sup><https://samate.nist.gov/SARD/test-suites/112>

<sup>18</sup>As of 4th Feb 2025

<sup>19</sup><https://nvd.nist.gov/developers/data-sources>

<sup>20</sup>As of 4th Feb 2025

TABLE XVI  
OVERVIEW OF SELECTED SE-WORKFLOW BENCHMARKS

Category	Benchmark	Language(s)	No. of Samples
Code Synthesis & Understanding	Methods2Test [212]	Java	780,944
	CRUXEval [213]	Python	800
	CRQBench [214]	C++	100
	CriticBench [215]	Python	3,825(across 5 tasks)
	CodeScope [216]	8 PLs	13,390 (across 8 tasks)
	CodeCriticBench [217]	Various	1,517 (Easy), 1,084 (Medium), 1,699 (Hard)
Merge Conflict Repair	CRUXEval-X [218]	19 PLs	19K
	ConflictBench [219]	Java	180
Type Inference	TypeEvalPy [220]	Python	845 (annotated labels)
	TypeEvalPy AutoGen [220]	Python	78373 (annotated labels)
Automatic Code Quality Review	CodeReview [221]	8 languages	7.9M pull requests
	Software Maintainability [222]	Java	519 projects
	BenMark [223]	Java	(evaluations of quality) 1,299,186 methods
	CodeReviewer [224]	9 languages	13,100
	CodeReview-New [225]	16 languages	14,600
	CodeReviewQA [226]	9 languages	900
Hallucination Detection	HALLUCODE [227]	Python	5,663
	CodeHaluEval [228]	Python	699
	Collu-Bench [229]	N/A	13,234
	LMDefects [230]	Java	60 (Easy), 53 (Medium)
	Codemirage [231]	Python	1,137

TABLE XVII  
OVERVIEW OF MULTI-CATEGORY BENCHMARKS, COVERING VARIOUS TASKS

Name	Language(s)	Tasks	Information
Big-Bench [233]	Python, Numeric, JSON, English	Functions over numbers, Mathematical Reasoning, Text2Code, Code2Text, Code explanation, Debugging, Turing Complete Concept Learning, amongst other tasks	250, several per category, 42, 60, 66, 34, 6,390
XLCoST [234]	C, C++, C#, Java, JS, Kotlin, PHP, Python, Ruby, Rust	Text2Code (program synthesis, code search), Code Summarization, Code Translation	567K (509k, 58k), 567K, 122K
CrossCodeBench [235]	Java, C#, Python, C++, JS, PHP, Go, Ruby, TS, C, Bash, Shell	Classification, In-Filling, Translation, Generation, Summarization, Type Prediction, Question Answering	6.6M, 13.4M, 2.4M, 19.5M, 11.2M, 773K, 190K
Long Code Arena [236]	English, Python, Java, Kotlin	Commit Message Generation, Module Summarization, Library-Based Code Generation, Project-Level Code Completion, Bug Localization, CI Builds Repair	163, 216, 150, 908 (varying sizes), 14.96K, 78
CodeXGLUE [236] MicrosoftDocs <sup>21</sup> CodeSearchNet [237]	English, Chinese, Norwegian, Danish, Latvian Go, Java, JS, PHP, Python, Ruby	Code Documentation Translation, Code Documentation (Code Summarization, Comment Generation)	(CN: 52K, NOR: 26K, DK: 45K, LT: 21K), 621870
DomainEval [238]	Python	Computation, Network, Basic operation, System, Visualization, Cryptography	5,892 cases total
CodeApex [23]	C++, English, Chinese	Programming comprehension Code generation, Code correction	250, 476, 1330

the existence of incorrect tests, suboptimal canonical solutions, and imprecise problem definitions amongst others.

### A. Approach

To improve the quality of the given benchmark, we pursue the following approach, also illustrated in Fig. 6. First, we initiate a comprehensive code review, leading to standardized observations (subsubsection V-A1). Then, we address these identified issues through a series of modifications (subsubsection V-A2), followed by a peer review (subsubsection V-A3) to ensure accuracy and reliability. Finally, we experiment with the revised




Fig. 6. BENCHFRAME's approach through a case-study of HUMANEVAL.

benchmark to evaluate and discuss the results. Below, more details are provided for specific steps in this approach.

1) *Standardized Observations in Current HumanEval Benchmarks:* Upon examining and manually reconstructing

TABLE XVIII  
OVERVIEW OF SELECTED BENCHMARKS FOR LOG STATEMENT  
GENERATION AND PARSING

Category	Benchmark	Language(s)	#Problems
Log Statement Generation	LANCE [239]	Java	76,421
	LogBench [240]	Java	6,849
	SCLogger [241]	Java	31,170
	AL-Bench [242]	Java	39,600
Log Parsing	LogBase [243]	Various	85,300
	LogHub [244]	Various	32,000
	LogHub-2.0 [245]	Various	3.6M
	LogPM [246]	Various	10,821,589
	LogEval [247]	Eng., Chin.	4K logs (4 tasks)

canonical solutions and experimenting with various HumanEval benchmark test suites, we identified several recurring issues. The system frequently produces incorrect and sub-optimal code, as canonical solutions are inefficient and fail to address critical assumptions outlined in the problem descriptions.

Additionally, these solutions often lack type annotations, further complicating the evaluation process. Another significant problem is the absence of quality testing. Test suites tend to overlap with example tests from the prompt, allowing incorrect canonical solutions to pass. Moreover, there are instances where the expected outputs in the test suites do not align with the canonical solutions' actual performance.

Compounding these issues is the poor quality of the problem descriptions, which contain grammatical errors, ambiguous instructions, and inconsistent formatting, particularly in the test examples. Furthermore, the system's support for language conversion frameworks, such as MultiPL-E, is inadequate. MultiPL-E, while the most comprehensive framework available, only supports equality assertions, which proves incompatible with the setup of many problems, further hindering the system's effectiveness.

2) *Modifications in HUMANEVALNEXT*: In HUMANEVALNEXT, we address all the above issues by manually modifying all problems in the original HumanEval benchmark. We decide to improve the original Human Eval as (1) flaws in the original version persist even in improved versions such as HumanEvalPlus, and (2) the fact that HumanEval is still widely used in new literature. While in our we have gone for the original benchmark, due to the adaptable nature of BenchFrame one could opt to have an improved version of the dataset as the starting point. To summarize, these are the general changes made in HUMANEVALNEXT and their benefits: In this work, several key improvements have been made to address the shortcomings of previous benchmarks. First, **all suboptimal and incorrect canonical solutions have been fixed**, which were previously missed due to insufficient testing and a lack of comprehensive quality review. Furthermore, **type annotations have been added** to all problems, offering valuable context and simplifying the translation to other programming languages.

The original HumanEval benchmark only included type annotations for the first 30 problems, representing merely 18% of the total set of 164 problems. In addition, **better support for language conversion frameworks** has been incorporated. For instance, HUMANEVALNEXT now features improved compatibility with frameworks like MultiPL-E [6], which supports translation to 18 additional programming languages by standardizing all tests to equality assertions where feasible. This change has reduced the number of incompatible problems by a factor of ten.

Besides these adjustments, challenging scenarios (such as negative values, zero instances, empty inputs, and non-alphanumeric symbols) are incorporated into each task to guarantee that only top-quality AI4SE models, adept at addressing diverse situations, succeed. Accordingly, **assertions are implemented** within the code wherever constraints are detailed in the problem description. This measure prevents models from ignoring crucial details, thus improving the benchmark's worth. In particular, we employ specification-based testing to assess functions; using boundary analysis, we explore combinations of within, on, and outside points. Furthermore, the **test examples in the problem descriptions have been refined**, which significantly affects model performance [15]. Problems with excessive test examples now feature a reduced set, distributing the evaluation workload more fairly. Lastly, **spelling errors have been corrected**, descriptions have been consistently formatted, and **problem descriptions have been aligned with the implementations**, while still leaving room for models to demonstrate intuitive problem-solving skills expected of high-quality AI4SE models. The difficulty level has also been raised by incorporating **more edge cases and modifying various problems**, aiming to better reflect real-world challenges faced by engineers and to mitigate issues related to data leakage and saturated performance on the HumanEval leaderboards. To illustrate the modifications of the tests in HumanEvalNext, consider Table XXI.

3) *Peer Review Process of HUMANEVALNEXT*: To ensure the accuracy and reliability of HUMANEVALNEXT, an independent reviewer verified all changes. This thorough review involved verifying the clarity and completeness of the problem descriptions, checking for consistency between the problem descriptions and the canonical solutions, and ensuring that both the solutions and test cases were correct and efficient. Where inefficiencies were identified, suggestions for optimization were provided. The review also scrutinized the test cases to ensure comprehensive coverage, identifying any gaps that could allow incorrect solutions to pass. For the sake of completeness, we must note that the peer review process is exclusively concerned with evaluating the modifications made and does not encompass other elements of the procedure. While the initial creation of the benchmark took over 100 hours, the independent peer-review process required an additional 16 hours. As a result of this review, 1% of the problems were redesigned due to structural issues, 9% received additional test cases, and 15% underwent minor grammar or clarity improvements. All suggested changes

<sup>21</sup><https://github.com/MicrosoftDocs/>

TABLE XIX  
COMMON LIMITATIONS IN AI4SE BENCHMARKS

Limitation	Description	Representative Examples	Frequency Observed
Language Specific	Focuses only on one programming language	HumanEval, MBPP, APPS	169
No Leaderboard	No official leaderboard; hinders fair and consistent comparisons	SPoC, RepoEval, CONCODE	201
Poor Maintenance/Data Quality	Dataset contains outdated problems, lacks documentation, or is not updated regularly	ToolQA, ExampleCheck	218
Not Peer Reviewed	Not been published in a peer-reviewed venue	CodeSearchNet, Methods2Test	79
Infrequently Used	Rarely cited in academic work	JavaBench, CoIR	199
Dataset Creation Method	How problems were obtained: <i>mined</i> from real-world sources, <i>created</i> by humans, or <i>LLM-generated</i>	AVATAR (mined), Spider (created), CRUXEval (generated)	Mined: 149, Created: 69, Generated: 46

TABLE XX  
QUESTIONNAIRE DESIGN OVERVIEW

Section	Questions	Scale	Additional Info
Participant Background	What is your professional background? What is your role? How familiar are you with AI4SE benchmarks? How many years of experience do you have in this field?	5-point Likert Scale (1: Not familiar, 5: Very familiar)	Experience question (< 1, 1–3, 3–5, 5+ years)
Search Functionality	How easy was it to navigate the interface? How intuitive was the search functionality? How effective was the tool in finding benchmarks? Was the visual evaluation of datasets useful?	5-point Likert Scale (1: Not useful, 5: Extremely useful)	N/A
Cross-referencing Feature	How useful was the cross-referencing feature? Did the tool help in understanding relationships between benchmarks? Was the visual interface for benchmark similarity useful?	5-point Likert Scale, Open-ended	Includes qualitative feedback option
User Experience & Feedback	How would you rate the overall user experience? How likely are you to use the tool in your work? Did you experience any issues or challenges? What other tools do you use for searching benchmarks? How does this tool compare to others? How well does the tool meet the needs of professionals in AI4SE?	Likert Scale, Open-ended	Open-ended questions for in-depth feedback

TABLE XXI  
COMPARISON OF TEST STATISTICS BETWEEN HUMAN EVAL (BASED ON <https://Github.com/Openai/Human-Eval/Blob/Master/Data/HumanEval.jsonl.gz> `human-Eval-v2-20210705.json`) AND HUMAN EVAL NEXT

Metric	HumanEval (original)	HumanEvalNext	$\Delta$
Total number of asserts	1325	2551	$\times 1.92$
Avg. number of asserts	8	16	$\times 2$
Med. number of asserts	7	11	$\times 1.57$
Min. number of asserts	1	4	+3
Total with < 5 asserts	34	2	-94%

were documented and reviewed by the original author, with every recommendation either implemented or refined further upon discussion. Since the peer-review involved modifications beyond the test suites, completions for all models were re-run. Despite these changes, the results remained largely consistent with the original, with 40% of models showing no change in their pass@1 scores, 50% showing a 1-2% absolute change, and only 10%—previously top performers—experiencing a 5% drop. This demonstrated that the peer-review process upheld the benchmark’s robustness while refining its quality.

### B. Experimental Setup

To assess the impact of the modifications applied to the benchmark, we examine the pass@1 performance of ten state-of-the-art open-source software code models using

the original HumanEval alongside two enhanced variants, HUMAN EVAL NEXT and EvalPlus [10]. We selected the top-performing models from the big code LLM leaderboards at the evaluation’s start: “NTQAI/Nxcode-CQ-7B-orpo”, “Qwen/CodeQwen1.5-7B”, “deepseek-ai/deepseek-coder-6.7b-instruct”, “TechxGenus/starcoder2-15b-instruct”, “ise-uiuc/Magicoder-S-DS-6.7B”, “Artigenz/Artigenz-Coder-DS-6.7B”, “HuggingFaceH4/starchat2-15b-v0.1”, “google/codegemma-7b-it”, “codeLlama/CodeLlama-13b-Instruct-hf”, and “Stabilityai/stable-code-3b”. For each task in the benchmark (164 total), the LLM is prompted using an instructional preamble asking the model to finish the implementation of the function requested in addition to providing the imports, function header, and function description with each request. We run the inference for the models on a cluster with one NVIDIA A100 80GB GPU and 32 CPU cores. During test execution, a timeout limit of 15 seconds is utilized per function call to disregard completions that are potentially looping forever or are considered overly inefficient with regard to the canonical solutions. Each test is executed in an evaluation suite using the precautions deployed by OpenAI.

### C. Results

Upon conducting the experiments outlined in subsection V-B, a key observation is an **average decrease of 31.22%** and a

TABLE XXII

PERFORMANCE COMPARISON ON HUMANEval BENCHMARKS (PASS@1) - VALUES IN PARENTHESES ARE THE Δ WITH THE BASELINE HUMANEval

Model	textbfHumanEval	HumanEval+	HumanEvalNext
stable-code	30.72	25.60 (-5.12)	1.83 (-28.89)
CodeLlama	50.60	34.10 (-16.50)	29.88 (-20.72)
codegemma	60.40	51.80 (-8.60)	41.46 (-18.94)
starchat2	73.80	71.30 (-2.50)	43.29 (-30.51)
Artigenz-Coder	75.60	72.60 (-3.00)	53.66 (-21.94)
Magocoder	76.80	71.30 (-5.50)	53.66 (-23.14)
starcoder2	77.40	60.00 (-17.40)	43.29 (-34.11)
deepseek-coder	80.22	71.30 (-8.92)	58.54 (-21.68)
CodeQwen1.5	87.20	45.70 (-41.50)	10.98 (-76.22)
Nxcode-CQ	87.23	83.50 (-3.73)	51.22 (-36.01)

TABLE XXIII

PERFORMANCE COMPARISON ON HUMANEval BENCHMARKS (PASS@1) - VALUES IN PARENTHESES ARE THE Δ WITH THE BASELINE HUMANEval

Model	HumanEval	HumanEval+	HumanEvalNext
stable-code	30.72	25.60 (-5.12)	1.83 (-28.89)
CodeLlama	50.60	34.10 (-16.50)	29.88 (-20.72)
codegemma	60.40	51.80 (-8.60)	41.46 (-18.94)
starchat2	73.80	71.30 (-2.50)	43.29 (-30.51)
Artigenz-Coder	75.60	72.60 (-3.00)	53.66 (-21.94)
Magocoder	76.80	71.30 (-5.50)	53.66 (-23.14)
starcoder2	77.40	60.00 (-17.40)	43.29 (-34.11)
deepseek-instruct	80.22	71.30 (-8.92)	58.54 (-21.68)
CodeQwen1.5	87.20	45.70 (-41.50)	10.98 (-76.22)
Nxcode-CQ	87.23	83.50 (-3.73)	51.22 (-36.01)




Fig. 7. Boxplot depicting the distribution of absolute drops in pass@1 score between HumanEval and the newly introduced HUMANEvalNext benchmark, based on 10 LLMs (Table XXIII).

**median decrease of 26.02%** (both absolute percentages) in pass@ 1 results of HumanEval when contrasted with the newly introduced HUMANEvalNext benchmark, using 10 different LLMs. Specific model outcomes are detailed in Table XXIII, with an overall depiction of the performance decreases displayed in Fig. 7. While there are still significant performance declines when comparing HumanEvalPlus with the original HumanEval (11.28 mean and 7.05 median), the declines are substantially larger with HumanEvalNext than with EvalPlus. Overall, the findings highlight a marked reduction in model performance when measured against HumanEvalNext as opposed to the initial HumanEval benchmark. This pattern emphasizes the enhanced difficulty and refined evaluation precision offered by HumanEvalNext, which incorporates more resilient environments featuring type annotations and clearer instructions.

A closer look reveals that the top-performing models from the original HumanEval benchmark do not maintain their standings in HUMANEvalNext. For example, while HUMANEvalNext consistently ranks deepseek-ai’s deepseek-coder-6.7b-instruct as the top performer, previous leaders like NTQAI’s Nxcode-CQ-7B-orpo and Qwen’s CodeQwen1.5-7B show a significant drop




Fig. 8. The flow-chart indicating the pipeline used to improve MBPP.

in their rankings. Specifically, NTQAI’s Nxcode-CQ-7B-orpo falls from an impressive 87.23% pass@1 in HumanEval to 51.22% in HUMANEvalNext, and Qwen’s CodeQwen1.5-7B plummets from 87.2% to 10.98%. This sharp decline suggests that certain models may have benefited from data leakage or other issues in the original HumanEval benchmark, indicating the necessity for a more rigorous benchmark like HUMANEvalNext to accurately assess model performance.

Especially the resilience of models, e.g., deepseek-ai/deepseek-coder-6.7b-instruct, can be confirmed by evaluating the model on the new challenges presented in HUMANEvalNext. When models also score relatively well in this benchmark, it highlights the models’ ability to adapt and perform competently under more demanding conditions, making HUMANEvalNext a great benchmark to reveal the reliability of the capabilities of models. This finding also emphasizes the need for regular updates to benchmarks, as reliance on outdated benchmarks can potentially misrepresent model capabilities over time, even when models claim not to train on the test data.

Furthermore, an interesting observation emerges when analyzing model size and performance: bigger is not always better. Larger models such as TechxGenus/starcoder2-15b-instruct and HuggingFaceH4/starchat2-15b-v0.1 do not consistently outperform smaller models. This observation suggests that model size alone is not a definitive predictor of success in complex, edge-case-inclusive benchmarks like HUMANEvalNext. For instance, looking at pass@1 scores, Techx-Genus/starcoder2-15b-instruct scores 77.4% on HumanEval but drops to 43.29% on HUMANEvalNext, while ise-uiuc/Magocoder-S-DS-6.7B scores 76.8% on HumanEval (lower) but only drops to 53.66% on HUMANEvalNext (higher). This highlights that increased model size does not necessarily equate to better performance in more challenging assessments.

Lastly, ranking problems based on their difficulty using pass metrics per problem, confirms that HUMANEvalNext presents a well-distributed range of complexity over the complete set of challenges. It also shows the increased difficulty of the benchmark, where even the easiest problems are not universally passed, with roughly 30% of the models still failing to solve them. Altogether, this distribution highlights HUMANEvalNext’s effectiveness in providing a thorough assessment environment, making it an excellent tool for evaluating straightforward coding capabilities of LLMs in a lightweight manner, the main reason behind the popularity of the original HumanEval benchmark and its variants.

## VI. DISCUSSION

In this section, we discuss the broader implications of our study; here, we attempt to tie our findings back to the RQs that guided our study. We, specifically, touch on what our results mean for researchers navigating the field, address the practical salability of our proposed approach, and outline threats to validity and opportunities for future work.

### A. Implications

1) *Finding the Right Benchmark*: In our analysis, and specifically RQ1 we find that AI4SE benchmarking is highly fragmented, with many benchmarks suffering from poor maintenance and lack of discoverability. With this finding, we underscore the fact a major challenge in AI4SE is the difficulty of selecting a contextually relevant benchmark. BenchScout was our direct answer to RQ2. With BenchScout, users can select more relevant benchmarks and gain deeper insights into model performance for their specific needs.

2) *Benchmarking the Benchmark*: Our answer to RQ3, BenchFrame, can have an impact on how we measure progress in AI4SE. Our results show the significant impact of integrating the BENCHFRAME approach and highlight the value of peer-reviewed, validated benchmarks. With many state-of-the-art models tested on similar benchmarks, both industry and academia should adapt their evaluation methods to ensure robust results. The refinement of benchmarks across various AI4SE tasks will be critical for guiding future research and ensuring that these models can perform effectively. Although this study, specifically in RQ1, highlights the issue of data leakage in current benchmarks, it remains true that HUMANEVALNEXT would also be affected by this issue. However, we argue that benchmarks should not be used indefinitely and should evolve to pose increasing challenges in alignment with model improvements.

### B. A Case Study on the agentification of BenchFrame

A potential criticism of the approach proposed in RQ3 is the significant manual effort it demands compared to a more automated system. To address these concerns about the generalizability and manual labor intensity of BenchFrame, we developed an agentic pipeline designed to replicate and automate the benchmark improvement process with minimal human intervention. Our baseline for this effort was the HUMANEVALNEXT dataset, which had previously been improved by manual human supervision. We aimed to determine whether an agentic approach could yield comparable improvements without the extensive human labor (to show noninferiority).

The agent pipeline was structured into three distinct phases, which mirrors the manual process described in subsection V-A2. In the first phase, the agent improved the problem description, which includes the function signature and the docstring. Using this refined description, the agent then generated an improved canonical solution. In the final phase, the agent created and set up the test cases to validate the solution. This process was applied across all the problems in the dataset,

and our hypothesis was that this automated approach would be noninferior to the human-improved dataset.

To validate this, we conducted a paired evaluation comparing the human- and agent-improved versions of each benchmark problem in the HUMANEVALNEXT dataset, specifically HUMANEVALNEXT versus HUMANEVALNEXT-Agentic. Two independent reviewers rated each pair using a five-point ordinal scale ranging from -2 (strongly preferring the human version) to +2 (strongly preferring the agentic version), with 0 indicating no preference. We tested whether the agentic pipeline was noninferior to the human-improved process, using a conservative noninferiority margin of  $\delta = -0.5$ . That is, we assessed whether the average rating for the agentic outputs was not more than 0.5 points worse than the human outputs.

For each reviewer, as well as for the combined ratings (averaged across reviewers), we performed both a one-sided one-sample t-test and a Wilcoxon signed-rank test to assess noninferiority. The results were consistent and conclusive: the mean ratings were 0.16 (SD = 0.67) for Reviewer 1, 0.53 (SD = 0.51) for Reviewer 2, and 0.35 (SD = 0.49) for the combined average. The corresponding t-statistics were 12.68, 25.75, and 22.10, with one-sided p-values all equal to 1.000 (indicating extremely strong evidence in favor of noninferiority given the direction of the test). The non-parametric Wilcoxon signed-rank tests further confirmed these results, with statistics of 12,155.5 ( $p = 3.46 \times 10^{-20}$ ), 13,198.0 ( $p = 7.20 \times 10^{-28}$ ), and 12,698.0 ( $p = 5.94 \times 10^{-27}$ ) respectively.

These results indicate that the agentic pipeline is noninferior to the human-improved process—indeed, with mean and median ratings above zero, the agentic output was often preferred. This supports the feasibility of using an agentic approach to automate benchmark refinement while significantly reducing manual effort.

While noninferior, we noted several standardized observations, in terms of pitfalls, during this evaluation, such as (1) instances where the agent's docstring occasionally revealed aspects of the solution, or (2) where canonical solutions made unintended assumptions, or (3) where the test cases did not adhere to the required, single-line assert, format. These observations were later accounted for in later stages of the feasibility study to make adjustments in the pipeline, and to ensure that it remained flexible and effective.

For the sake of completeness, we present the refined pipeline used in the process of benchmark refinement. The pipeline begins by initializing an attempt counter, which tracks the number of iterations taken to complete a given benchmark improvement. The process then proceeds through three core phases: (a) Text Improvement, in which the agent refines the problem description and adds a type-annotated function signature; (b) Code Improvement, where the canonical implementation is modified according to Python best practices and based on the improved task description; and (c) Test Improvement, which involves generating assert-based test cases that include edge cases, boundary conditions, and common error scenarios. These three components are then assembled into a complete Python program and passed to the Validation phase (d), which runs the implementation against the test suite in a sandboxed subprocess.

TABLE XXIV  
PASS@1 RESULTS FOR HUMANEVAL, HUMANEVALNEXT, AND  
HUMANEVALNEXT-AGENTIC

Model	HumanEval	HumanEvalNext- Agentic ( $\Delta$ )	HumanEvalNext ( $\Delta$ )
stable-code	30.72	2.44(-28.28)	1.83 (-28.89)
CodeLlama	50.60	18.90(-31.70)	29.88 (-20.72)
codegemma	60.40	22.56 (-37.84)	41.46 (-18.94)
starchat2	73.80	26.83 (-46.97)	43.29 (-30.51)
Artigenz-Coder	75.60	31.70 (-43.90)	53.66 (-21.94)
Magicoder	76.80	31.70 (-45.10)	53.66 (-23.14)
starcoder2	77.40	14.02 (-63.38)	43.29 (-34.11)
deepseek-coder	80.22	35.37 (-44.85)	58.54 (-21.68)
CodeQwen1.5	87.20	1.22 (-86.00)	10.98 (-76.22)
Nxcode-CQ	87.23	31.70 (-55.53)	51.22 (-36.01)

If all tests pass on the first try, the process terminates successfully. Otherwise, the pipeline evaluates whether the number of attempts is below a fixed threshold (three, in our case). If so, the pipeline either increments the attempt counter and retries the full process or, alternatively, in the case that more tests are passing than were present in the original test suite, the pipeline identifies a minimal set of passing tests (e) to be used as the new test suite. This loop ensures that even partially successful test runs can be used to improve the benchmark iteratively. After three failed attempts, the task is marked as unsuccessful.

For our experiments, we used the OpenAI API<sup>22</sup>, specifically, the `o3-mini-2025-01-31` as the underlying model for the agent. We evaluated this new dataset on the same set of models, and the results (as presented in Table XXIV) indicated comparable (and even better) performance patterns. These findings provide adequate justification for how an agentic pipeline can effectively generalize and provide a semi-automated solution for benchmark improvement.

The results presented indicate, that while a certain level of human intervention, specifically in the review phase, is still required to ensure quality of the improvements, it is but a fraction of the total time required to manually improve the benchmark. From a cost perspective, we can also report that the total incurred cost from calling the apis for the models was \$5.28<sup>23</sup>. This is but a small portion of the labor costs required to improve the benchmark manually. The introduction of this agentic layer significantly reduced the manual effort required. We claim that, with this approach, one can extend it to other datasets with similar success.

### C. Evaluating the Generalizability of BenchFrame

An additional critique that can be raised regarding our approach is its generalizability. To tackle this issue, we adopted a strategy akin to HumanEval. Using our systematic observations of MBPP from the review, we apply BenchFrame to 100 problems selected from a pool of 500 MBPP test samples. The rationale behind this selection is twofold: firstly, it aims to assess a new task, specifically program synthesis, and secondly, it considers the benchmark's popularity. During this process,

<sup>22</sup><https://platform.openai.com>

<sup>23</sup>At the time of running these experiments, on the 31st of July 2025

TABLE XXV  
PASS@1 RESULTS FOR MBPP AND MBPPNEXT.  
BASED ON <https://Github.com/Google-Research/Google-Research/Blob/Master/Mbpp/Mbpp.json>  
ACCESSED ON 14-07-2025

Model	MBPP	MBPPNext ( $\Delta$ )
stable-code	34.0	28.0 (-6%)
CodeLlama	40.0	31.0 (-9%)
codegemma	43.0	33.0 (-10%)
starchat2	59.0	42.0 (-17%)
Artigenz-Coder	60.0	47.0 (-13%)
Magicoder	63.0	45.0 (-18%)
starcoder2	44.0	36.0 (-8%)
deepseek-coder	68.0	50.0 (-18%)
CodeQwen1.5	53.0	40.0 (-13%)
Nxcode-CQ	74.0	52.0 (-22%)

we subject the function descriptions, implementations (including variable typing), and significantly extend test cases; more specifically, we go from a fixed 3 tests per problem to an average of 12.35, with a median of 11.5. We finally go through several rounds of peer review to ensure the quality of the improvements. We refer to the resulting benchmark as MBPPNext. The evaluation results of this dataset in comparison to the baseline are detailed in Table XXV. We observe a comparable decline in performance similar to that noted for HumanEval, specifically, an average decrease of 13.4 percentage points across the methods evaluated. These findings further underscore the necessity for thorough peer review and verification of benchmarks, as discussed in previous sections of this paper.

### D. Future Work

Future work on BENCHFRAME will focus on expanding to additional programming languages, which it already has been optimized for, yet these variants have not been produced or evaluated. Although evaluating ten LLMs provides valuable insight, it remains unclear how larger, top-performing models behind paywalls, such as GPT and Gemini, would perform; Future research could evaluate the performance differences of such models when comparing the base and the pro version. Furthermore, future research should focus on applying the underlying ideas of BenchFrame to mutli-file and project-level benchmarks like Defects4J.

### E. Threats to the Validity

**Construct Validity** : To reduce bias and errors in the literature review, two authors followed a structured protocol for selecting and filtering sources. A threat to construct validity comes from the subjectivity in defining and applying the inclusion/exclusion criteria. To address this, we pre-defined clear criteria and peer-reviewed the selection process for consistency. Another potential validity threat stems from the design of our user study, as it may not fully capture the tool's overall functionality, strengths, or weaknesses. To mitigate this, we evaluated diverse aspects including usability, functionality, usefulness, and intuitiveness. Lastly, to minimize biases during HUMANEVALNEXT's development, we applied consistent

refinement criteria and peer-reviews, although subjective interpretation remains a residual risk.

**Internal Validity:** A potential threat in the user study is selection bias. To mitigate this, we included participants from both industry and academia, ensuring a range of skills and experience. To reduce the real-world impact arising from any study biases, we have created a continuous feedback mechanism in BenchScout to receive continuous feedback from users to be able to meet their needs. All participants received the same tool, guidance, and instructions. To reduce bias in the peer review or HUMANEVALNEXT, the reviewer was not informed of specific changes made by the first author.

**External Validity:** A threat here is the generalizability of the user study's results. We mitigated this by including 22 participants, but the sample size may still limit the generalizability of the results. Although we evaluated the effects of the BENCHFRAME approach on the performance of ten models with one of the most widely used benchmarks, this could not be considered sufficient to prove the generalizability of the results; adding more models and application to more benchmarks could further confirm our results. BenchFrame may be criticized for its limited applicability due to the significant manual effort it demands compared to more automated systems like EvalPlus. We demonstrate that the process of improving benchmarks such as HumanEval and MBPP can be (semi-)automated through the use of agentic pipelines.

## VII. RELATED WORK

The rapid integration of LLMs into software engineering has led to a corresponding proliferation of benchmarks designed to evaluate them. This has spurred research to systematically map this new terrain, assess the quality of evaluation resources, and build tools for navigating them. Our work is situated within these three emerging areas.

*a) AI4SE Surveys and Taxonomies:* Initial research in this area included broad surveys of LLMs in software engineering, which established evaluation as a central research challenge [174], [248], [249], [250]. More recently, meta-analyses have focused on the benchmarks themselves. For instance, Wang et al. created a taxonomy based on the Software Development Life Cycle, revealing a research gap with a strong focus on code implementation and a negligence towards design and requirements engineering [251]. Our work contributes to this area by providing a comprehensive review of 273 benchmarks from 247 studies (RQ1). We go beyond categorization by analyzing each benchmark across a 14-point metadata schema to identify systemic limitations, such as poor maintenance, language specificity, and a lack of peer review.

*b) Tools for Benchmark Discovery and Navigation:* While platforms like Hugging Face and the now-discontinued Papers With Code serve as valuable repositories for hosting datasets [252], they often lack the specialized search and visualization capabilities needed by AI4SE researchers. BenchScout addresses this gap (RQ2) by providing an extensible semantic search tool specifically designed for the AI4SE community.

Unlike general-purpose repositories, it combines semantic embeddings with structured metadata filters and an interactive visualization of the benchmark landscape to improve the discoverability and selection of relevant evaluation tools.

*c) Methodologies for Improving Benchmark Quality:* There is a growing agreement that many static benchmarks are prone to saturation and data contamination, which can compromise the validity of evaluation results [18]. In response, two movements have arisen. Efforts like LiveCodeBench fall under the “build anew” category; these strive to develop dynamic, “live” benchmarks that continuously gather new challenges from real-world environments to avoid overfitting.

Our method, BenchFrame, belongs to the “repair and refine” category and focuses on improving the quality of existing benchmarks (RQ3). A method within this category is DyCodeEval, which uses seed contexts to frame problems in specific scenarios (e.g., banking, healthcare, education). This approach is designed to produce more reliable results by minimizing data contamination and memorization effects [253]. Our approach sets itself apart by incorporating a peer-review-oriented process that includes standardized observations, precise modifications, and independent validation. Like DyCodeEval, we show its practical scalability through an agentic pipeline that automates the improvement process and tackles the main challenge, namely, the manual efforts required in BenchFrame.

## VIII. CONCLUSION

The findings of our study highlight the importance of reliable and consistent benchmarking in AI4SE to drive the development of more robust models. Through the creation of BenchFrame and the enhancement of the HumanEval benchmark, we have demonstrated that higher-quality benchmarks reveal substantial performance gaps, as shown by the 31.2% average reduction in pass@1 scores across ten state-of-the-art models. This significant decline highlights the impact of more stringent evaluations. BenchScout further enhances this process by facilitating the discovery of relevant benchmarks and reduces the overhead associated with selecting the appropriate tools for evaluation.

## DATA AVAILABILITY

We publicly release the results of our literature review, user study, and 50% of the manually refined benchmark.<sup>24</sup> Upon acceptance, the complete benchmark will be made available on both GitHub and HuggingFace.

## REFERENCES

- [1] M. Chen et al., “Evaluating large language models trained on code,” 2021. [Online]. Available: <https://arxiv.org/abs/2107.03374>
- [2] G. G. Team, “Gemini: A family of highly capable multimodal models,” 2023. [Online]. Available: [https://storage.googleapis.com/deepmind-media/gemini/gemini\\_1\\_report.pdf](https://storage.googleapis.com/deepmind-media/gemini/gemini_1_report.pdf)
- [3] OpenAI, “Gpt-4 technical report,” 2023. [Online]. Available: <https://arxiv.org/abs/2303.08774>
- [4] T. van Dam, F. van der Heijden, P. de Bekker, B. Nieuwschepen, M. Otten, and M. Izadi, “Investigating the performance of language models

<sup>24</sup><https://github.com/AISE-TUdelft/AI4SE-benchmarks>

- for completing code in functional programming languages: A Haskell case study,” 2024, [Online]. Available: <https://arxiv.org/abs/2403.15185>
- [5] B. Athiwaratkun et al., “Multi-lingual evaluation of code generation models,” 2022. [Online]. Available: <https://arxiv.org/abs/2210.14868>
  - [6] F. Cassano et al., “MultiPL-E: A scalable and extensible approach to benchmarking neural code generation,” 2022. [Online]. Available: <https://arxiv.org/abs/2208.08227>
  - [7] N. Muennighoff et al., “OctoPACK: Instruction tuning code large language models,” 2023. [Online]. Available: <https://arxiv.org/abs/2308.07124>
  - [8] Q. Zheng et al., “CodeGeeX: A pre-trained model for code generation with multilingual evaluations on HumanEval-X,” 2023. [Online]. Available: <http://arxiv.org/abs/2303.17568>
  - [9] Y. Dong, J. Ding, X. Jiang, G. Li, Z. Li, and Z. Jin, “CodeScore: Evaluating code generation by learning code execution,” 2023. [Online]. Available: <https://arxiv.org/abs/2301.09043>
  - [10] 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,” 2023. [Online]. Available: <https://arxiv.org/abs/2305.01210>
  - [11] I. CODAIT, “AI for code: Predict code complexity using IBM’S CodeNet dataset,” 2021. [Online]. Available: <https://community.ibm.com/community/user/datascience/blogs/sepideh-seifzadeh1/2021/10/05/ai-for-code-predict-code-complexity-using-ibms-cod>
  - [12] Q. Peng, Y. Chai, and X. Li, “HumanEval-XL: A multilingual code generation benchmark for cross-lingual natural language generalization,” Mar. 2024. [Online]. Available: <http://arxiv.org/abs/2402.16694>
  - [13] N. Raihan, A. Anastasopoulos, and M. Zampieri, “mHumanEval—A multilingual benchmark to evaluate large language models for code generation,” 2024, *arXiv:2410.15037*.
  - [14] C. S. Xia, Y. Deng, and L. Zhang, “Top leaderboard ranking = Top coding proficiency, always? EvoEval: Evolving coding benchmarks via LLM,” 2024. [Online]. Available: <https://arxiv.org/abs/2403.19114>
  - [15] J. Austin et al., “Program synthesis with large language models,” 2021. [Online]. Available: <https://arxiv.org/abs/2108.07732>
  - [16] Y. Li et al., “Competition-level code generation with AlphaCode,” *Science*, vol. 378, no. 6624, pp. 1092–1097, 2022, doi: 10.1126/science.abq1158.
  - [17] D. Hendrycks et al., “Measuring coding challenge competence with APPS,” 2021. [Online]. Available: <https://arxiv.org/abs/2105.09938v3>
  - [18] N. Jain et al., “LiveCodeBench: Holistic and contamination free evaluation of large language models for code,” Jun. 2024. [Online]. Available: <http://arxiv.org/abs/2403.07974>
  - [19] H. Tian et al., “Is ChatGPT the ultimate programming assistant – How far is it?” Aug. 2023. [Online]. Available: <http://arxiv.org/abs/2304.11938>
  - [20] S. Quan et al., “CodeElo: Benchmarking competition-level code generation of LLMs with human-comparable Elo ratings,” Jan. 2025. [Online]. Available: <http://arxiv.org/abs/2501.01257>
  - [21] Q. Dougherty and R. Mehta, “Proving the coding interview: A benchmark for formally verified code generation,” 2025, *arXiv:2502.05714*.
  - [22] Z. Xu, Y. Liu, Y. Yin, M. Zhou, and R. Poovendran, “KodCode: A diverse, challenging, and verifiable synthetic dataset for coding,” 2025, *arXiv:2503.02951*.
  - [23] Z. Liu, X. Hu, D. Zhou, L. Li, X. Zhang, and Y. Xiang, “Code generation from flowcharts with texts: A benchmark dataset and an approach,” in *Proc. Findings Assoc. Comput. Linguistics (EMNLP)*, 2022, pp. 6069–6077.
  - [24] Z. Wang, S. Liu, Y. Sun, H. Li, and K. Shen, “CodeContests+: High-quality test case generation for competitive programming,” 2025, *arXiv:2506.05817*.
  - [25] Z. Zheng et al., “LiveCodeBench Pro: How do Olympiad medalists judge LLMs in competitive programming?” 2025, *arXiv:2506.11928*.
  - [26] J. Sikka, K. Satya, Y. Kumar, S. Uppal, R. R. Shah, and R. Zimmermann, “Learning based methods for code runtime complexity prediction,” 2019. [Online]. Available: <https://arxiv.org/abs/1911.01155>
  - [27] K. Moudgalya, A. Ramakrishnan, V. Chemudupati, and X. H. Lu, “Tasty: A transformer based approach to space and time complexity,” 2023. [Online]. Available: <https://arxiv.org/abs/2305.05379>
  - [28] S.-Y. Baik, M. Jeon, J. Hahn, J. Kim, Y.-S. Han, and S.-K. Ko, “CodeComplex: A time-complexity dataset for bilingual source codes,” 2024. [Online]. Available: <https://arxiv.org/abs/2401.08719>
  - [29] A. Yadav and M. Singh, “PythonSaga: Redefining the benchmark to evaluate code generating LLM,” 2024. [Online]. Available: <https://arxiv.org/abs/2401.03855v2>
  - [30] D. Huang, J. M. Zhang, Y. Qing, and H. Cui, “EffiBench: Benchmarking the efficiency of automatically generated code,” 2024. [Online]. Available: <https://arxiv.org/abs/2402.02037v2>
  - [31] M. Weyssow, A. Kamanda, and H. Sahraoui, “CodeUltraFeedback: An LLM-as-a-judge dataset for aligning large language models to coding preferences,” 2024, *arXiv:2403.09032*.
  - [32] A. Shypula et al., “Learning performance-improving code edits,” Apr. 2024. [Online]. Available: <http://arxiv.org/abs/2302.07867>
  - [33] Y. Peng, J. Wan, Y. Li, and X. Ren, “Coffe: A code efficiency benchmark for code generation,” *Proc. ACM Softw. Eng.*, vol. 2, no. FSE, pp. 242–265, 2025.
  - [34] Y. Lai et al., “DS-1000: A natural and reliable benchmark for data science code generation,” 2022. [Online]. Available: <https://arxiv.org/abs/2211.11501>
  - [35] D. Zan et al., “CERT: Continual pre-training on sketches for library-oriented code generation,” 2022. [Online]. Available: <https://arxiv.org/abs/2206.06888>
  - [36] R. Agashe, S. Iyer, and L. Zettlemoyer, “JuICE: A large scale distantly supervised dataset for open domain context-based code generation,” in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP-IJCNLP)*, China: Assoc. for Comput. Linguistics, 2019, pp. 5436–5446. [Online]. Available: <https://aclanthology.org/D19-1546>
  - [37] S. Chandel, C. B. Clement, G. Serrato, and N. Sundaresan, “Training and evaluating a jupyter notebook data science assistant,” 2022. [Online]. Available: <https://arxiv.org/abs/2201.12901>
  - [38] J. Huang et al., “Execution-based evaluation for data science code generation models,” 2022, *arXiv:2211.09374*.
  - [39] Y. Zhang, Q. Jiang, X. Han, N. Chen, Y. Yang, and K. Ren, “Benchmarking data science agents,” Feb. 2024. [Online]. Available: <http://arxiv.org/abs/2402.17168>
  - [40] D. Zan, B. Chen, Z. Lin, B. Guan, W. Yongji, and J.-G. Lou, “When language model meets private library,” in *Proc. Findings Assoc. Comput. Linguistics (EMNLP)*, 2022, pp. 277–288.
  - [41] X. Tang, B. Qian, R. Gao, J. Chen, X. Chen, and M. B. Gerstein, “BioCoder: A benchmark for bioinformatics code generation with large language models,” *Bioinformatics*, vol. 40, no. S1, pp. i266–i276, Jul. 2024.
  - [42] Y. Cui, “WebApp1K: A practical code-generation benchmark for web app development,” 2024. [Online]. Available: <https://arxiv.org/abs/2408.00019>
  - [43] D. Hendrycks et al., “Measuring coding challenge competence with APPS,” in *Proc. NeurIPS Datasets Benchmarks*, 2021.
  - [44] H. Lightman et al., “Let’s verify step by step,” May 2023. [Online]. Available: <http://arxiv.org/abs/2305.20050>
  - [45] A. Amini, S. Gabriel, P. Lin, R. Koncel-Kedziorski, Y. Choi, and H. Hajishirzi, “MathQA: Towards interpretable math word problem solving with operation-based formalisms,” 2019. [Online]. Available: <https://arxiv.org/abs/1905.13319>
  - [46] S. Mishra et al., “Lila: A unified benchmark for mathematical reasoning,” 2022, *arXiv:2210.17517*.
  - [47] S. Roy and D. Roth, “Solving general arithmetic word problems,” in *Proc. Conf. Empirical Methods Natural Lang. Process.*, Lisbon, Portugal: Assoc. for Comput. Linguistics, 2015, pp. 1743–1752. [Online]. Available: <https://aclanthology.org/D15-1202>
  - [48] L. Gao et al., “PAL: Program-aided language models,” 2022. [Online]. Available: <https://arxiv.org/abs/2211.10435>
  - [49] W. Chen et al., “TheoremQA: A theorem-driven question answering dataset,” in *Proc. Conf. Empirical Methods Natural Lang. Process.*, Singapore: Assoc. for Comput. Linguistics, 2023, pp. 7889–7901. [Online]. Available: <https://aclanthology.org/2023.emnlp-main.489>
  - [50] P. Haller, J. Golde, and A. Akbik, “PECC: Problem extraction and coding challenges,” in *Proc. Int. Conf. Lang. Resour. Eval.*, 2024, pp. 12690–12699.
  - [51] H. Su et al., “BRIGHT: A realistic and challenging benchmark for reasoning-intensive retrieval,” 2024. [Online]. Available: <https://arxiv.org/abs/2407.12883>
  - [52] P. Yin, B. Deng, E. Chen, B. Vasilescu, and G. Neubig, “Learning to mine aligned code and natural language pairs from stack overflow,” in *Proc. 15th Int. Conf. Mining Softw. Repositories (MSR)*, New York, NY, USA: ACM, 2018, pp. 476–486, doi: 10.1145/3196398.3196408.
  - [53] Z. Wang, G. Cuenca, S. Zhou, F. F. Xu, and G. Neubig, “MCoNaLa: A benchmark for code generation from multiple natural languages,” 2022. [Online]. Available: <https://arxiv.org/abs/2203.08388>
  - [54] G. Orlanski and A. Gittens, “Reading StackOverflow encourages cheating: Adding question text improves extractive code generation,” Jun. 2021. [Online]. Available: <http://arxiv.org/abs/2106.04447>

- [55] Y. Hao et al., "AixBench: A code generation benchmark dataset," 2022. [Online]. Available: <https://arxiv.org/abs/2206.13179>
- [56] J. Huang et al., "CoSQA: 20,000+ web queries for code search and question answering," in *Proc. 59th Annu. Meeting Assoc. Comput. Linguistics 11th Int. Joint Conf. Natural Lang. Process.*, ACL, 2021, pp. 5690–5700. [Online]. Available: <https://aclanthology.org/2021.acl-long.442>
- [57] S. Lu et al., "CodeXGLUE: A machine learning benchmark dataset for code understanding and generation," 2021. [Online]. Available: <http://arxiv.org/abs/2102.04664>
- [58] S. Iyer, I. Konstas, A. Cheung, and L. Zettlemoyer, "Mapping language to code in programmatic context," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, Assoc. for Comput. Linguistics, 2018, pp. 1643–1652, doi: 10.18653/v1/D18-1192.
- [59] E. Nijkamp et al., "CodeGen: An open large language model for code with multi-turn program synthesis," 2023. [Online]. Available: <https://arxiv.org/abs/2203.13474>
- [60] S. Zhang, J. Wang, G. Dong, J. Sun, Y. Zhang, and G. Pu, "Experimenting a new programming practice with LLMs," 2024. [Online]. Available: <https://arxiv.org/abs/2401.01062>
- [61] P. Liguori et al., "Can we generate shellcodes via natural language? An empirical study," *Autom. Softw. Eng.*, vol. 29, no. 1, p. 30, 2022.
- [62] T. Helmuth, T. Helmuth, P. Kelly, and P. Kelly, "PSB2: The second program synthesis benchmark suite," in *Proc. Annu. Conf. Genetic Evol. Comput.*, 2021, pp. 785–794.
- [63] R. Li et al., "TACO: Topics in algorithmic COde generation dataset," 2023, *arXiv:2312.14852*.
- [64] S. Honarvar, M. V. D. Wilk, and A. Donaldson, "Turbulence: Systematically and automatically testing instruction-tuned large language models for code," 2023, *arXiv:2312.14856*.
- [65] J. Shin, M. Wei, J. Wang, L. Shi, and S. Wang, "The good, the bad, and the missing: Neural code generation for machine learning tasks," *ACM Trans. Softw. Eng. Methodol.*, vol. 33, no. 2, pp. 1–24, Feb. 2024. [Online]. Available: <https://dl.acm.org/doi/10.1145/3630009>
- [66] J. Chen et al., "RMCBench: Benchmarking large language models' resistance to malicious code," in *Proc. 39th IEEE/ACM Int. Conf. Autom. Softw. Eng.*, Oct. 2024, pp. 995–1006. [Online]. Available: <http://arxiv.org/abs/2409.15154>
- [67] P. Liguori et al., "EVIL: Exploiting software via natural language," in *Proc. IEEE 32nd Int. Symp. Softw. Rel. Eng. (ISSRE)*, Oct. 2021, pp. 321–332. [Online]. Available: <http://arxiv.org/abs/2109.00279>
- [68] Y. Xie, A. Xie, D. Sheth, P. Liu, D. Fried, and C. Rose, "CodeBenchGen: Creating scalable execution-based code generation benchmarks," 2024, *arXiv:2404.00566*.
- [69] K. Yan, H. Guo, X. Shi, J. Xu, Y. Gu, and Z. Li, "CodeIF: Benchmarking the instruction-following capabilities of large language models for code generation," 2025, *arXiv:2502.19166*.
- [70] P. Wang et al., "CodeIF-Bench: Evaluating instruction-following capabilities of large language models in interactive code generation," 2025, *arXiv:2503.22688*.
- [71] P. Yin et al., "Natural language to code generation in interactive data science notebooks," in *Proc. 61st Annu. Meeting Assoc. Comput. Linguistics (Volume 1: Long Papers)*, 2023, pp. 126–173.
- [72] N. Shah, Z. Genc, and D. Araci, "StackEval: Benchmarking LLMs in coding assistance," in *Proc. Adv. Neural Inf. Process. Syst.*, 2024, vol. 37, pp. 36976–36994.
- [73] I. Petrukha, Y. Kurliak, and N. Stulova, "SwiftEval: Developing a language-specific benchmark for LLM-generated code evaluation," in *Proc. IEEE/ACM 2nd Int. Conf. AI Found. Models Softw. Eng. (Forge)*, Piscataway, NJ, USA: IEEE Press, 2025, pp. 73–77.
- [74] J. Gong et al., "CoSQA+: Pioneering the multi-choice code search benchmark with test-driven agents," 2025. [Online]. Available: <https://arxiv.org/abs/2406.11589>
- [75] InfiCoder, "InfiCoder-Eval: Systematically evaluating question-answering for code large language models," 2023. [Online]. Available: <https://openreview.net/forum?id=wncrC9rmHXh>
- [76] X. Hu, G. Li, X. Xia, D. Lo, and Z. Jin, "Deep code comment generation," in *Proc. 26th Conf. Program Comprehension (ICPC)*, New York, NY, USA: ACM, 2018, pp. 200–210, doi: 10.1145/3196321.3196334.
- [77] X. Hu, G. Li, X. Xia, D. Lo, and Z. Jin, "Deep code comment generation with hybrid lexical and syntactical information," *Empirical Softw. Eng.*, vol. 25, no. 3, pp. 2179–2217, May 2020. [Online]. Available: <http://link.springer.com/10.1007/s10664-019-09730-9>
- [78] X. Jin, J. Larson, W. Yang, and Z. Lin, "Binary code summarization: Benchmarking ChatGPT/GPT-4 and other large language models," 2023. [Online]. Available: <https://arxiv.org/abs/2312.09601>
- [79] M. Allamanis, M. Allamanis, H. Peng, H. Peng, C. Sutton, and C. Sutton, "A convolutional attention network for extreme summarization of source code," 2016, *arXiv:1602.03001*.
- [80] A. LeClair, A. LeClair, S. Jiang, S. Jiang, C. McMillan, and C. McMillan, "A neural model for generating natural language summaries of program subroutines," in *Proc. Int. Conf. Softw. Eng.*, 2019, pp. 795–806.
- [81] X. Hu, G. Li, X. Xia, D. Lo, S. Lu, and Z. Jin, "Summarizing source code with transferred API knowledge," in *Proc. 27th Int. Joint Con. Artif. Intell.*, 2018, pp. 2269–2275.
- [82] M. Hasan et al., "CoDesc: A large code-description parallel dataset," May 2021. [Online]. Available: <http://arxiv.org/abs/2105.14220>.
- [83] A. V. M. Barone and R. Sennrich, "A parallel corpus of Python functions and documentation strings for automated code documentation and code generation." [Online]. Available: <https://aclanthology.org/I17-2053/>
- [84] K. Pai, P. Devanbu, and T. Ahmed, "CoDocBench: A dataset for code-documentation alignment in software maintenance," 2025. [Online]. Available: <https://arxiv.org/abs/2502.00519>
- [85] C. Hu, Y. Chai, H. Zhou, F. Meng, J. Zhou, and X. Gu, "How effectively do code language models understand poor-readability code?" in *Proc. 39th IEEE/ACM Int. Conf. Autom. Softw. Eng.*, 2024, pp. 795–806.
- [86] S. Yun, S. Lin, X. Gu, and B. Shen, "Project-specific code summarization with in-context learning," *J. Syst. Softw.*, vol. 216, Oct. 2024, Art. no. 112149.
- [87] J. Li et al., "Can LLM already serve as a database interface? A big bench for large-scale database grounded text-to-SQLs," in *Proc. Neural Inf. Process. Syst.*, 2023, 42330–42357.
- [88] C.-H. Lee, O. Polozov, and M. Richardson, "KaggleDBQA: Realistic evaluation of text-to-SQL parsers," in *Proc. Annu. Meeting Assoc. Comput. Linguistics*, 2021.
- [89] Z. Yao et al., "StaQC: A systematically mined question-code dataset from stack overflow," in *Proc. Web Conf.*, 2018, pp. 1693–1703.
- [90] F. Lei et al., "Spider 2.0: Evaluating language models on real-world enterprise text-to-SQL workflows," 2024. [Online]. Available: <https://arxiv.org/abs/2411.07763>
- [91] Y. Gan et al., "Towards robustness of text-to-SQL models against synonym substitution," Jun. 2021. [Online]. Available: <http://arxiv.org/abs/2106.01065>
- [92] X. Deng, A. H. Awadallah, C. Meek, O. Polozov, H. Sun, and M. Richardson, "Structure-grounded pretraining for text-to-SQL," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Human Lang. Technol.*, 2021, pp. 1337–1350. [Online]. Available: <http://arxiv.org/abs/2010.12773>
- [93] Y. Gan, X. Chen, and M. Purver, "Exploring underexplored limitations of cross-domain text-to-SQL generalization," Sep. 2021. [Online]. Available: <http://arxiv.org/abs/2109.05157>
- [94] Q. Min, Y. Shi, and Y. Zhang, "A pilot study for Chinese SQL semantic parsing," Oct. 2019. [Online]. Available: <http://arxiv.org/abs/1909.13293>
- [95] T. Yu et al., "SParC: Cross-domain semantic parsing in context," Jun. 2019. [Online]. Available: <http://arxiv.org/abs/1906.02285>
- [96] Q. Liang et al., "Lyra: A benchmark for turducken-style code generation," in *Proc. 31st Int. Joint Conf. Artif. Intell.*, Jul. 2022, pp. 4238–4244. [Online]. Available: <http://arxiv.org/abs/2108.12144>
- [97] L. Wang et al., "DuSQL: A large-scale and pragmatic Chinese text-to-SQL dataset," in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, Association for Computational Linguistics, 2020, pp. 6923–6935, [Online]. Available: <https://www.aclweb.org/anthology/2020.emnlp-main.562>
- [98] T. Yu et al., "CoSQL: A conversational text-to-SQL Challenge towards cross-domain natural language interfaces to databases," Sep. 2019. [Online]. Available: <http://arxiv.org/abs/1909.05378>
- [99] H. Li et al., "OmniSQL: Synthesizing high-quality text-to-SQL data at scale," 2025, *arXiv:2503.02240*.
- [100] D. Bakshandaeva, O. Somov, E. Dmitrieva, V. Davydova, and E. Tutubalina, "PAUQ: Text-to-SQL in Russian," in *Proc. Findings Assoc. Comput. Linguistics (EMNLP)*, Y. Goldberg, Z. Kozareva, and Y. Zhang, Eds. Abu Dhabi, United Arab Emirates: Assoc. for Comput. Linguistics, Dec. 2022, pp. 2355–2376. [Online]. Available: <https://aclanthology.org/2022.findings-emnlp.175/>
- [101] S. Almohaimeed, S. Almohaimeed, M. Al Ghanim, and L. Wang, "Ar-spider: Text-to-SQL in Arabic," in *Proc. 39th ACM/SIGAPP Symp. Appl. Comput. (SAC)*, New York, NY, USA: Association for Computing Machinery, 2024, pp. 1024–1030, doi: 10.1145/3605098.3636065.

- [102] A. B. Kanburoğlu and F. B. Tek, “TUR2SQL: A cross-domain Turkish dataset for text-to-SQL,” in *Proc. 8th Int. Conf. Comput. Sci. Eng. (UBMK)*, 2023, pp. 206–211.
- [103] K. L. Aw, S. Montariol, B. Alkhamissi, M. Schrimpf, and A. Bosselut, “Instruction-tuning aligns LLMs to the human brain,” 2023. [Online]. Available: <https://arxiv.org/abs/2312.00575>
- [104] B. Li et al., “DevBench: A comprehensive benchmark for software development,” 2024. [Online]. Available: <https://arxiv.org/abs/2403.08604>
- [105] J. Li et al., “DevEval: A manually-annotated code generation benchmark aligned with real-world code repositories,” 2024. [Online]. Available: <https://arxiv.org/abs/2405.19856>
- [106] Z. Zeng, Y. Wang, R. Xie, W. Ye, and S. Zhang, “CoderUJB: An executable and unified Java benchmark for practical programming scenarios,” 2024. [Online]. Available: <https://arxiv.org/abs/2403.19287>
- [107] Y. Zhuang, Y. Yu, K. Wang, H. Sun, and C. Zhang, “ToolQA: A dataset for LLM question answering with external tools,” Jun. 2023. [Online]. Available: <http://arxiv.org/abs/2306.13304>
- [108] X. Wang et al., “MINT: Evaluating LLMs in multi-turn interaction with tools and language feedback,” Mar. 2024. [Online]. Available: <http://arxiv.org/abs/2309.10691>
- [109] L. Gong, S. Wang, M. Elhoushi, and A. Cheung, “Evaluation of LLMs on syntax-aware code fill-in-the-middle tasks,” Jun. 2024. [Online]. Available: <http://arxiv.org/abs/2403.04814>
- [110] X. Liu, et al., “AgentBench: Evaluating LLMs as agents,” Oct. 2023. [Online]. Available: <http://arxiv.org/abs/2308.03688>
- [111] Y. Xiao, R. Wang, L. Kong, D. Golac, and W. Wang, “CSR-Bench: Benchmarking LLM agents in deployment of computer science research repositories,” 2025, *arXiv:2502.06111*.
- [112] X. Du, et al., “ClassEval: A manually-crafted benchmark for evaluating LLMs on class-level code generation,” 2023. [Online]. Available: <https://arxiv.org/abs/2308.01861>
- [113] T. Y. Zhuo et al., “BigCodeBench: Benchmarking code generation with diverse function calls and complex instructions,” 2024. [Online]. Available: <https://arxiv.org/abs/2406.15877>
- [114] S. Wang, L. Ding, L. Shen, Y. Luo, B. Du, and D. Tao, “OOP: Object-oriented programming evaluation benchmark for large language models,” in *Proc. Findings Assoc. Comput. Linguistics (ACL)*, 2024, pp. 13619–13639.
- [115] M. K. Roy et al., “CodeSense: A real-world benchmark and dataset for code semantic reasoning,” 2025. [Online]. Available: <https://arxiv.org/abs/2506.00750>
- [116] P. XueLi et al., “ClassEval-T: Evaluating large language models in class-level code translation,” *Proc. ACM Softw. Eng.*, vol. 2, no. ISSSTA, Jun. 2025, pp. 1421–1444, doi: 10.1145/3728940.
- [117] C. E. Jimenez et al., “SWE-Bench: Can language models resolve real-world GitHub issues?” 2023. [Online]. Available: <https://arxiv.org/abs/2310.06770>
- [118] Y. Ding et al., “CrossCodeEval: A diverse and multilingual benchmark for cross-file code completion,” 2023. [Online]. Available: <https://arxiv.org/abs/2310.11248>
- [119] H. Yu et al., “CoderEval: A benchmark of pragmatic code generation with generative pre-trained models,” 2023. [Online]. Available: <https://arxiv.org/abs/2302.00288>
- [120] L. A. Agrawal, A. Kanade, N. Goyal, S. K. Lahiri, and S. K. Rajamani, “Guiding language models of code with global context using monitors,” Nov. 2023. [Online]. Available: <http://arxiv.org/abs/2306.10763>
- [121] J. Svajlenko and C. K. Roy, “Evaluating clone detection tools with BigCloneBench,” in *Proc. IEEE Int. Conf. Softw. Maintenance Evol. (ICSM/E)*, Bremen, Germany: IEEE Press, Sep. 2015, pp. 131–140. [Online]. Available: <http://ieeexplore.ieee.org/document/7332459/>
- [122] L. Zhang et al., “DI-BENCH: Benchmarking large language models on dependency inference with testable repositories at scale,” 2025. [Online]. Available: <https://arxiv.org/abs/2501.13699>
- [123] I. Bouzenia, B. P. Krishan, and M. Pradel, “DyPyBench: A benchmark of executable Python software,” *Proc. ACM Softw. Eng.*, vol. 1, no. FSE, pp. 338–358, Jul. 2024, doi: 10.1145/3643742.
- [124] D. Zan et al., “Multi-SWE-Bench: A multilingual benchmark for issue resolving,” 2025, *arXiv:2504.02605*.
- [125] A. Ouyang et al., “KernelBench: Can LLMs Write Efficient GPU Kernels?” 2025, *arXiv:2502.10517*.
- [126] K. Cheng et al., “CODEMENV: Benchmarking large language models on code migration,” 2025, *arXiv:2506.00894*.
- [127] J. Guo et al., “CodeEditorBench: Evaluating code editing capability of LLMs,” in *Proc. ICLR 3rd Workshop Deep Learn. Code*, 2025.
- [128] K. Liu et al., “ProjectEval: A benchmark for programming agents automated evaluation on project-level code generation,” 2025, *arXiv:2503.07010*.
- [129] T. Liu, C. Xu, and J. McAuley, “RepoBench: Benchmarking repository-level code auto-completion systems,” 2023. [Online]. Available: <https://arxiv.org/abs/2306.03091>
- [130] F. Zhang et al., “RepoCoder: Repository-level code completion through iterative retrieval and generation,” 2023. [Online]. Available: <https://arxiv.org/abs/2303.12570>
- [131] J. Li, G. Li, X. Zhang, Y. Dong, and Z. Jin, “EvoCodeBench: An evolving code generation benchmark aligned with real-world code repositories,” 2024. [Online]. Available: <https://arxiv.org/abs/2404.00599>
- [132] D. Zan et al., “Codes: Natural language to code repository via multi-layer sketch,” 2024, *arXiv:2403.16443*.
- [133] X. Tang et al., “ML-Bench: Evaluating large language models and agents for machine learning tasks on repository-level code,” Aug. 2024. [Online]. Available: <http://arxiv.org/abs/2311.09835>
- [134] M. Liu, T. Yang, Y. Lou, X. Du, Y. Wang, and X. Peng, “Code-Gen4Libs: A two-stage approach for library-oriented code generation,” in *Proc. 38th IEEE/ACM Int. Conf. Autom. Softw. Eng. (ASE)*, Luxembourg, Luxembourg: IEEE Press, Sep. 2023, pp. 434–445. [Online]. Available: <https://ieeexplore.ieee.org/document/10298327/>
- [135] I. Badertdinov et al., “SWE-Rebench: An automated pipeline for task collection and decontaminated evaluation of software engineering agents,” 2025, *arXiv:2505.20411*.
- [136] M. S. Rashid et al., “SWE-PolyBench: A multi-language benchmark for repository level evaluation of coding agents,” 2025, *arXiv:2504.08703*.
- [137] D. Zheng et al., “HumanEvo: An evolution-aware benchmark for more realistic evaluation of repository-level code generation,” in *proc. IEEE/ACM 47th Int. Conf. Softw. Eng. (ICSE)*, Washington, DC, USA: IEEE Computer Society, 2025, pp. 764–764.
- [138] S. Liang, Y. Hu, N. Jiang, and L. Tan, “Can language models replace programmers? REPOCOD Says ‘not yet,’” 2024, *arXiv:2410.21647*.
- [139] W. Li et al., “FEA-Bench: A benchmark for evaluating repository-level code generation for feature implementation,” 2025, *arXiv:2503.06680*.
- [140] J. Cao, Z. Chen, J. Wu, S.-C. Cheung, and C. Xu, “JavaBench: A benchmark of object-oriented code generation for evaluating large language models,” in *Proc. 39th IEEE/ACM Int. Conf. Automat. Softw. Eng.*, 2024, pp. 870–882.
- [141] L. Zhang et al., “SWE-bench goes live!” 2025, *arXiv:2505.23419*.
- [142] S. Miserendino, M. Wang, T. Patwardhan, and J. Heidecke, “SWE-Lancer: Can frontier LLMs earn one million from real-world freelance software engineering?” 2025, *arXiv:2502.12115*.
- [143] Y. Song et al., “RestGPT: Connecting large language models with real-world restful APIs,” 2023. [Online]. Available: <https://arxiv.org/abs/2306.06624>
- [144] Y. Peng et al., “Revisiting, benchmarking and exploring API recommendation: How far are we?” 2021. [Online]. Available: <https://arxiv.org/abs/2112.12653>
- [145] Q. Huang, X. Xia, Z. Xing, D. Lo, and X. Wang, “API method recommendation without worrying about the task-API knowledge gap,” in *Proc. 33rd ACM/IEEE Int. Conf. Autom. Softw. Eng. (ASE)*, New York, NY, USA: ACM, 2018, pp. 293–304, doi: 10.1145/3238147.3238191.
- [146] S. G. Patil, T. Zhang, X. Wang, and J. E. Gonzalez, “Gorilla: Large language model connected with massive APIs,” 2023. [Online]. Available: <https://arxiv.org/abs/2305.15334>
- [147] M. Li et al., “API-bank: A comprehensive benchmark for tool-augmented LLMs,” 2023. [Online]. Available: <https://arxiv.org/abs/2304.08244>
- [148] Z. Z. Wang et al., “CodeRAG-Bench: Can retrieval augment code generation?” 2024. [Online]. Available: <https://arxiv.org/abs/2406.14497>
- [149] N. Rao, C. Bansal, and J. Guan, “Search4Code: Code search intent classification using weak supervision,” Mar. 2021. [Online]. Available: <http://arxiv.org/abs/2011.11950>
- [150] X. Li, et al., “CoIR: A comprehensive benchmark for code information retrieval models,” Jul. 2024. [Online]. Available: <http://arxiv.org/abs/2407.02883>
- [151] A. Al-Kaswan, M. Izadi, and A. Van Deursen, “Traces of memorisation in large language models for code,” in *Proc. IEEE/ACM 46th Int. Conf. Softw. Eng.*, Lisbon, Portugal: ACM, Apr. 2024, pp. 1–12. [Online]. Available: <https://dl.acm.org/doi/10.1145/3597503.3639133>
- [152] T. Zhang, G. Upadhyaya, A. Reinhardt, H. Rajan, and M. Kim, “Are code examples on an online Q&A forum reliable?: A study of API misuse on stack overflow,” in *Proc. IEEE/ACM 40th Int. Conf. Softw. Eng. (ICSE)*, New York, NY, USA: IEEE Press, 2018, pp. 886–896.
- [153] L. Zhong and Z. Wang, “Can LLM replace stack overflow? A study on robustness and reliability of large language model code generation,”

- in *Proc. AAAI Conf. Artif. Intell.*, vol. 38, no. 19, pp. 21841–21849, 2024.
- [154] Z. Gu, J. Wu, J. Liu, M. Zhou, and M. Gu, “An empirical study on API misuse bugs in open-source c programs,” in *Proc. IEEE 43rd Annu. Comput. Softw. Appl. Conf. (COMPSAC)*, Piscataway, NJ, USA: IEEE Press, 2019, vol. 1, pp. 11–20.
- [155] S. Kulal et al., “SPOC: Search-based pseudocode to code,” in *Proc. Adv. Neural Inf. Process. Syst.*, Curran Associates, 2019, vol. 32, pp. 11906–11917. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2019/file/7298332f04ac004a0ca444c69ecf6f6b-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/7298332f04ac004a0ca444c69ecf6f6b-Paper.pdf)
- [156] M. Zavershynskiy, A. Skidanov, and I. Polosukhin, “NAPS: Natural program synthesis dataset,” 2018. [Online]. Available: <https://arxiv.org/abs/1807.03168>
- [157] J. Wu, S. Chen, J. Cao, H. C. Lo, and S.-C. Cheung, “Isolating language-coding from problem-solving: Benchmarking LLMs with PseudoEval,” 2025. [Online]. Available: <https://arxiv.org/abs/2502.19149>
- [158] Y. Oda et al., “Learning to generate pseudo-code from source code using statistical machine translation,” in *Proc. 30th IEEE/ACM Int. Conf. Autom. Softw. Eng. (ASE)*, 2015, pp. 574–584.
- [159] V. Zhong, C. Xiong, and R. Socher, “Seq2SQL: Generating structured queries from natural language using reinforcement learning,” 2017. [Online]. Available: <https://arxiv.org/abs/1709.00103>
- [160] T. Yu et al., “Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task,” 2018. [Online]. Available: <https://arxiv.org/abs/1809.08887>
- [161] X. V. Lin et al., “NL2Bash: A corpus and semantic parser for natural language interface to the Linux operating system,” in *Proc. Computat. Lang.*, 2018.
- [162] B. Roziere, J. M. Zhang, F. Charton, M. Harman, G. Synnaeve, and G. Lample, “Leveraging automated unit tests for unsupervised code translation,” 2022. [Online]. Available: <https://arxiv.org/abs/2110.06773>
- [163] M. Zhu, K. Suresh, and C. K. Reddy, “Multilingual code snippets training for program translation,” in *Proc. AAAI Conf. Artif. Intell.*, 2022, vol. 36, no. 10, pp. 11783–11790, doi: 10.1609/aaai.v36i10.21434.
- [164] W. U. Ahmad, M. G. R. Tushar, S. Chakraborty, and K.-W. Chang, “Avatar: A parallel corpus for Java-Python program translation,” in *Proc. Findings Assoc. Computat. Linguistics*, Toronto, Canada: Assoc. for Comput. Linguistics, 2023, pp. 2268–2281. [Online]. Available: <https://aclanthology.org/2023.findings-acl.143>
- [165] W. Yan, Y. Tian, Y. Li, Q. Chen, and W. Wang, “CodeTransOcean: A comprehensive multilingual benchmark for code translation,” 2023. [Online]. Available: <https://arxiv.org/abs/2310.04951>
- [166] M. Jiao, T. Yu, X. Li, G. Qiu, X. Gu, and B. Shen, “On the evaluation of neural code translation: Taxonomy and benchmark,” in *Proc. 38th IEEE/ACM Int. Conf. Automated Softw. Eng. (ASE)*, Luxembourg, Luxembourg: IEEE Press, Sep. 2023, pp. 1529–1541. [Online]. Available: <https://ieeexplore.ieee.org/document/10298408/>
- [167] J. Zhang, P. Nie, J. J. Li, and M. Gligoric, “Multilingual code co-evolution using large language models,” in *Proc. 31st ACM Joint Eur. Softw. Eng. Conf. Symp. Found. Softw. Eng.*, San Francisco CA USA: ACM, Nov. 2023, pp. 695–707. [Online]. Available: <https://dl.acm.org/doi/10.1145/3611643.3616350>
- [168] G. Ou, M. Liu, Y. Chen, X. Peng, and Z. Zheng, “Repository-level code translation benchmark targeting rust,” 2024, *arXiv:2411.13990*.
- [169] P. Jana, P. Jha, H. Ju, G. Kishore, A. Mahajan, and V. Ganesh, “CoTran: An LLM-based code translator using reinforcement learning with feedback from compiler and symbolic execution,” in *Proc. ECAI*, IOS Press, 2024, pp. 4011–4018.
- [170] Y. Wang et al., “RepoTransBench: A real-world benchmark for repository-level code translation,” 2024, *arXiv:2412.17744*.
- [171] I. Paul, G. Glavaš, and I. Gurevych, “IRCoder: Intermediate representations make language models robust multilingual code generators,” Apr. 2024. [Online]. Available: <http://arxiv.org/abs/2403.03894>
- [172] R. Just, D. Jalali, and M. D. Ernst, “Defects4J: A database of existing faults to enable controlled testing studies for java programs,” in *Proc. Int. Symp. Softw. Testing Anal. (ISSTA)*, New York, NY, USA: ACM, 2014, pp. 437–440, doi: 10.1145/2610384.2628055.
- [173] A. Silva, N. Saavedra, and M. Monperrus, “GitBug-Java: A reproducible benchmark of recent Java bugs,” in *Proc. IEEE/ACM 21st Int. Conf. Mining Software Repositories (MSR)*, 2024, pp. 118–122, doi: 10.48550/arXiv.2402.02961.
- [174] Q. Zhang et al., “A critical review of large language model on software engineering: An example from ChatGPT and automated program repair,” 2023. [Online]. Available: <https://arxiv.org/abs/2310.08879>
- [175] B. Yang et al., “CREF: An LLM-based conversational software repair framework for programming tutors,” 2024, *arXiv:2406.13972*.
- [176] J. Y. Lee, S. Kang, J. Yoon, and S. Yoo, “The GitHub recent bugs dataset for evaluating LLM-based debugging applications,” 2023, *arXiv:2310.13229*.
- [177] C. L. Goues et al., “The manybugs and introclass benchmarks for automated repair of c programs,” *IEEE Trans. Softw. Eng.*, vol. 41, no. 12, pp. 1236–1256, Dec. 2015.
- [178] R. Tian et al., “DebugBench: Evaluating debugging capability of large language models,” 2024, *arXiv:2401.04621*.
- [179] D. Lin et al., “QuixBugs: A multi-lingual program repair benchmark set based on the Quixey challenge,” in *Proc. ACM SIGPLAN Int. Conf. Syst., Program., Lang. Appl., Softw. Humanity*, 2017, pp. 55–56.
- [180] B. Labash, A. Rosedale, A. Reents, L. Negritto, and C. Wiel, “Res-Q: Evaluating code-editing large language model systems at the repository scale,” 2024, *arXiv:2406.16801*.
- [181] H. M. Babe, S. Nguyen, Y. Zi, A. Guha, M. Q. Feldman, and C. J. Anderson, “StudentEval: A benchmark of student-written prompts for large language models of code,” 2023, *arXiv:2306.04556*.
- [182] Y. Hu, U. Z. Ahmed, S. Mechtaev, B. Leong, and A. Roychoudhury, “Re-factoring based program repair applied to programming assignments,” in *Proc. 34th IEEE/ACM Int. Conf. Autom. Softw. Eng. (ASE)*, 2019, pp. 388–398.
- [183] Y. Wu, Z. Li, J. M. Zhang, and Y. Liu, “ConDefects: A new dataset to address the data leakage concern for LLM-based fault localization and program repair,” Oct. 2023. [Online]. Available: <http://arxiv.org/abs/2310.16253>
- [184] R. Sharifdeen, M. Mirchev, Y. Noller, and A. Roychoudhury, “Cerberus: A program repair framework,” in *Proc. IEEE/ACM 45th Int. Conf. Softw. Eng.: Companion Proc. (ICSE-Companion)*, Melbourne, Australia: IEEE Press, May 2023, pp. 73–77. [Online]. Available: <https://ieeexplore.ieee.org/document/10172676/>
- [185] A. Silva and M. Monperrus, “RepairBench: Leaderboard of frontier models for program repair,” in *Proc. IEEE/ACM Int. Workshop Large Lang. Models Code (LLM4Code)*, Piscataway, NJ, USA: IEEE Press, 2025, pp. 9–16.
- [186] H. Nunes, T. Sharma, and E. Figueiredo, “MaRV: A manually validated refactoring dataset,” in *Proc. IEEE/ACM 2nd Int. Conf. AI Found. Models Softw. Eng. (Forge)*, Piscataway, NJ, USA: IEEE Press, 2025, pp. 141–145.
- [187] R. Widiasari et al., “BugsInPy: A database of existing bugs in python programs to enable controlled testing and debugging studies,” in *Proc. 28th ACM Joint Meeting Eur. Softw. Eng. Conf. Symp. Found. Softw. Eng.*, 2020, pp. 1556–1560.
- [188] G. P. Bhandari, P. Bhandari, A. Naseer, A. Naseer, L. Moonen, and L. Moonen, “CVEfixes: Automated collection of vulnerabilities and their fixes from open-source software,” in *Proc. Softw. Eng.*, 2021, pp. 30–39.
- [189] C. Tony, M. Mutas, N. E. D. Ferreyra, and R. Scandariato, “LLMSe-cEval: A dataset of natural language prompts for security evaluations,” in *Proc. IEEE Work. Conf. Mining Softw. Repositories*, 2023, pp. 588–592.
- [190] M. L. Siddiq, M. L. Siddiq, J. C. S. Santos, and J. C. S. Santos, “SecurityEval dataset: Mining vulnerability examples to evaluate machine learning-based code generation techniques,” 2022. [Online]. Available: <https://doi.org/10.1109/MSR59073.2023.00084>
- [191] Q.-C. Bui, R. Scandariato, and N. E. D. Ferreyra, “Vul4J: A dataset of reproducible Java vulnerabilities geared towards the study of program repair techniques,” in *Proc. 19th Int. Conf. Mining Softw. Repositories*, Pittsburgh Pennsylvania: ACM, May 2022, pp. 464–468. [Online]. Available: <https://dl.acm.org/doi/10.1145/3524842.3528482>
- [192] N. Tihanyi, T. Bisztray, R. Jain, M. A. Ferrag, L. C. Cordeiro, and V. Mavroeidis, “The FormAI dataset: Generative AI in software security through the lens of formal verification,” in *Proc. 19th Int. Conf. Predictive Models Data Analytics Softw. Eng.*, San Francisco CA USA: ACM, Dec. 2023, pp. 33–43. [Online]. Available: <https://dl.acm.org/doi/10.1145/3617555.3617874>
- [193] Y. Wu et al., “How effective are neural networks for fixing security vulnerabilities,” in *Proc. 32nd ACM SIGSOFT Int. Symp. Softw. Testing Anal.*, Jul. 2023, pp. 1282–1294. [Online]. Available: <http://arxiv.org/abs/2305.18607>
- [194] T. Durieux, J. F. Ferreira, R. Abreu, and P. Cruz, “Empirical review of automated analysis tools on 47,587 Ethereum smart contracts,” in *Proc. ACM/IEEE 42nd Int. Conf. Softw. Eng.*, Jun. 2020, pp. 530–541. [Online]. Available: <http://arxiv.org/abs/1910.10601>
- [195] Y. Zhou, S. Liu, J. Siow, X. Du, and Y. Liu, “Devign: Effective vulnerability identification by learning comprehensive program semantics via graph neural networks,” Sep. 2019. [Online]. Available: <http://arxiv.org/abs/1909.03496>

- [196] Y. Zheng, et al., “D2A: A Dataset built for AI-Based vulnerability detection methods using differential analysis,” Feb. 2021. [Online]. Available: <http://arxiv.org/abs/2102.07995>
- [197] J. Fan, Y. Li, S. Wang, and T. N. Nguyen, “A C/C++ code vulnerability dataset with code changes and CVE summaries,” in *Proc. 17th Int. Conf. Mining Softw. Repositories (MSR)*, New York, NY, USA: ACM, 2020, pp. 508–512, doi: 10.1145/3379597.3387501.
- [198] X. Mei et al., “ARVO: Atlas of reproducible vulnerabilities for open source software,” 2024, *arXiv:2408.02153*.
- [199] E. T. Liu, A. Wang, S. Mateega, C. Georgescu, and D. Tang, “VADER: A human-evaluated benchmark for vulnerability assessment, detection, explanation, and remediation,” 2025, *arXiv:2505.19395*.
- [200] B. Lin, S. Wang, L. Chen, and X. Mao, “There are more fish in the sea: Automated vulnerability repair via binary templates,” 2024, *arXiv:2411.18088*.
- [201] M. Tufano, S. Chandel, A. Agarwal, N. Sundaresan, and C. B. Clement, “Predicting code coverage without execution,” 2023. [Online]. Available: <https://arxiv.org/abs/2307.13383>
- [202] C. Watson, M. Tufano, K. Moran, G. Bavota, and D. Shshyanyk, “On learning meaningful assert statements for unit test cases,” in *Proc. ACM/IEEE 42nd Int. Conf. Softw. Eng.*, Jun. 2020, pp. 1398–1409. [Online]. Available: <http://arxiv.org/abs/2002.05800>
- [203] Z. Wang, K. Liu, G. Li, and Z. Jin, “HITS: High-coverage LLM-based unit test generation via method slicing,” in *Proc. 39th IEEE/ACM Int. Conf. Autom. Softw. Eng.*, Sacramento, CA, USA: ACM, Oct. 2024, pp. 1258–1268. [Online]. Available: <https://dl.acm.org/doi/10.1145/3691620.3695501>
- [204] P. Bareiß, B. Souza, M. d’Amorim, and M. Pradel, “Code generation tools (almost) for free? A study of few-shot, pre-trained language models on code,” Jun. 2022. [Online]. Available: <http://arxiv.org/abs/2206.01335>
- [205] C. Wan et al., “Automated testing of software that uses machine learning APIs,” in *Proc. 44th Int. Conf. Softw. Eng.*, Pittsburgh, PA, USA: ACM, May 2022, pp. 212–224. [Online]. Available: <https://dl.acm.org/doi/10.1145/3510003.3510068>
- [206] Z. Zeng, Y. Wang, R. Xie, W. Ye, and S. Zhang, “CoderUJB: An executable and unified Java benchmark for practical programming scenarios,” in *Proc. 33rd ACM SIGSOFT Int. Symp. Softw. Testing Anal.*, 2024, pp. 124–136.
- [207] Q. Zhang, Y. Shang, C. Fang, S. Gu, J. Zhou, and Z. Chen, “Testbench: Evaluating class-level test case generation capability of large language models,” 2024, *arXiv:2409.17561*.
- [208] W. Wang et al., “TestEval: Benchmarking large language models for test case generation,” 2024, *arXiv:2406.04531*.
- [209] A. S. Yaraghi, D. Holden, N. Kahani, and L. Briand, “Automated test case repair using language models,” *IEEE Trans. Softw. Eng.*, vol. 51, no. 4, pp. 1104–1133, Apr. 2025.
- [210] Y. Wang et al., “ProjectTest: A project-level unit test generation benchmark and impact of error fixing mechanisms,” 2025, *arXiv:2502.06556*.
- [211] J. Xu, B. Pang, J. Qu, H. Hayashi, C. Xiong, and Y. Zhou, “Clover: A test case generation benchmark with coverage, long-context, and verification,” 2025. [Online]. Available: <https://arxiv.org/abs/2502.08806>
- [212] M. Tufano, D. Drain, A. Svyatkovskiy, S. K. Deng, and N. Sundaresan, “Unit test case generation with transformers and focal context,” 2020. [Online]. Available: <https://arxiv.org/abs/2009.05617>
- [213] A. Gu, B. Roziere, H. Leather, A. Solar-Lezama, G. Synnaeve, and S. I. Wang, “CruxEval: A benchmark for code reasoning, understanding and execution,” 2024. [Online]. Available: <https://arxiv.org/abs/2401.03065>
- [214] E. Dinella, S. Chandra, and P. Maniatis, “CRQBench: A benchmark of code reasoning questions,” 2024 [Online]. Available: <https://arxiv.org/abs/2408.08453>.
- [215] Z. Lin, Z. Gou, T. Liang, R. Luo, H. Liu, and Y. Yang, “CriticBench: Benchmarking LLMs for critique-correct reasoning,” Jun. 2024. [Online]. Available: <http://arxiv.org/abs/2402.14809>
- [216] W. Yan, et al., “CodeScope: An execution-based multilingual multitask multidimensional benchmark for evaluating LLMs on code understanding and generation,” Jun. 2024. [Online]. Available: <http://arxiv.org/abs/2311.08588>
- [217] A. Zhang et al., “CodeCriticBench: A holistic code critique benchmark for large language models,” 2025, *arXiv:2502.16614*.
- [218] R. Xu et al., “CruxEval-X: A benchmark for multilingual code reasoning, understanding and execution,” 2025. [Online]. Available: <https://arxiv.org/abs/2408.13001>
- [219] B. Shen and N. Meng, “ConflictBench: A benchmark to evaluate software merge tools,” *J. Syst. Softw.*, vol. 214, 2024, Art. no. 112084. [Online]. Available: <https://dl.acm.org/doi/10.1016/j.jss.2024.112084>
- [220] A. P. S. Venkatesh, S. Sabu, J. Wang, A. M. Mir, L. Li, and E. Bodden, “TypeEvalpy: A micro-benchmarking framework for python type inference tools,” in *Proc. IEEE/ACM 46th Int. Conf. Softw. Eng., Companion Proc. (ICSE-Companion)*, 2023, pp. 49–53.
- [221] Z. Li et al., “Automating code review activities by large-scale pre-training,” in *Proc. ESEC/SIGSOFT FSE*, 2022, pp. 1035–1047.
- [222] M. Schnappinger, M. Schnappinger, A. Fietzke, A. Fietzke, A. Pretschner, and A. Pretschner, “Defining a software maintainability dataset: Collecting, aggregating and analysing expert evaluations of software maintainability,” in *Proc. IEEE Int. Conf. Softw. Maintenance Evol.*, 2020, pp. 278–289.
- [223] T. Wang, Y. Zhang, L. Jiang, Y. Tang, G. Li, and H. Liu, “Deep learning based identification of inconsistent method names: How far are we?” *Empirical Softw. Eng.*, vol. 30, no. 1, 2025, Art. no. 31.
- [224] Z. Li et al., “Automating code review activities by large-scale pre-training,” in *Proc. 30th ACM Joint Eur. Softw. Eng. Conf. Symp. Found. Softw. Eng.*, 2022, pp. 1035–1047.
- [225] Q. Guo et al., “Exploring the potential of ChatGPT in automated code refinement: An empirical study,” in *Proc. 46th IEEE/ACM Int. Conf. Softw. Eng.*, 2024, pp. 1–13.
- [226] H. Y. Lin, C. Liu, H. Gao, P. Thongtanunam, and C. Treude, “CodeReviewQA: The code review comprehension assessment for large language models,” 2025, *arXiv:2503.16167*.
- [227] F. Liu et al., “Exploring and evaluating hallucinations in LLM-powered code generation,” 2024, *arXiv:2404.00971*.
- [228] Y. Tian et al., “CodeHalu: Investigating code hallucinations in LLMs via execution-based verification,” in *Proc. AAAI Conf. Artif. Intell.*, 2025, vol. 39, pp. 25300–25308.
- [229] N. Jiang, Q. Li, L. Tan, and T. Zhang, “Collu-Bench: A benchmark for predicting language model hallucinations in code,” 2024, *arXiv:2410.09997*.
- [230] Z. Fan, X. Gao, M. Mirchev, A. Roychoudhury, and S. H. Tan, “Automated repair of programs from large language models,” in *Proc. IEEE/ACM 45th Int. Conf. Softw. Eng. (ICSE)*, Piscataway, NJ, USA: IEEE Press, 2023, pp. 1469–1481.
- [231] V. Agarwal, Y. Pei, S. Alamir, and X. Liu, “CodeMirage: Hallucinations in code generated by large language models,” 2024, *arXiv:2408.08333*.
- [232] M. L. Siddiq, J. C. S. Santos, R. H. Tanvir, N. Ulfat, F. A. Rifat, and V. C. Lopes, “Using large language models to generate JUnit tests: An empirical study,” 2024. [Online]. Available: <https://arxiv.org/abs/2305.00418>
- [233] A. Srivastava et al., “Beyond the imitation game: Quantifying and extrapolating the capabilities of language models,” 2022. [Online]. Available: <https://arxiv.org/abs/2206.04615>
- [234] M. Zhu, A. Jain, K. Suresh, R. Ravindran, S. Tipirneni, and C. K. Reddy, “XLCoST: A benchmark dataset for cross-lingual code intelligence,” 2022. [Online]. Available: <https://arxiv.org/abs/2206.08474>
- [235] C. Niu, C. Li, V. Ng, and B. Luo, “CrossCodeBench: Benchmarking cross-task generalization of source code models,” in *Proc. Int. Conf. Softw. Eng.*, 2023, pp. 537–549.
- [236] E. Bogomolov, et al., “Long code arena: A set of benchmarks for long-context code models,” 2024. [Online]. Available: <https://arxiv.org/abs/2406.11612>
- [237] H. Husain, H.-H. Wu, T. Gazit, M. Allamanis, and M. Brockschmidt, “CodeSearchNet challenge: Evaluating the state of semantic code search,” 2019. [Online]. Available: <https://arxiv.org/abs/1909.09436>
- [238] Q. Zhu et al., “DomainEval: An auto-constructed benchmark for multi-domain code generation,” in *Proc. AAAI Conf. Artif. Intell.*, 2025, vol. 39, pp. 26148–26156.
- [239] A. Mastroaolo, L. Pascarella, and G. Bavota, “Using deep learning to generate complete log statements,” in *Proc. IEEE/ACM 44th Int. Conf. Softw. Eng. (ICSE)*, 2022, pp. 2279–2290.
- [240] Y. Li et al., “Exploring the effectiveness of LLMs in automated logging generation: An empirical study,” 2023, *arXiv:2307.05950*.
- [241] Y. Li et al., “Go static: Contextualized logging statement generation,” *Proc. ACM Softw. Eng.*, vol. 1, no. FSE, pp. 609–630, 2024.
- [242] B. Tan, J. Xu, Z. Zhu, and P. He, “AL-Bench: A benchmark for automatic logging,” 2025. [Online]. Available: <https://arxiv.org/abs/2502.03160>
- [243] C. Zhang et al., “LogBase: A large-scale benchmark for semantic log parsing,” *Proc. ACM Softw. Eng.*, vol. 2, no. ISSTA, pp. 2091–2112, 2025.

- [244] J. Zhu, S. He, P. He, J. Liu, and M. R. Lyu, "Loghub: A large collection of system log datasets for ai-driven log analytics," in *Proc. IEEE 34th Int. Symp. Softw. Rel. Eng. (ISSRE)*, Piscataway, NJ, USA: IEEE Press, 2023, pp. 355–366.
- [245] Z. Jiang et al., "A large-scale evaluation for log parsing techniques: How far are we?" in *Proc. 33rd ACM SIGSOFT Int. Symp. Softw. Testing Anal.*, 2024, pp. 223–234.
- [246] S. Hashemi, J. Nyssölä, and M. V. Mäntylä, "LogPM: Character-based log parser benchmark," in *Proc. IEEE Int. Conf. Softw. Anal., Evolution Reengineering (SANER)*, 2024, pp. 705–710.
- [247] T. Cui et al., "LogEval: A comprehensive benchmark suite for large language models in log analysis," 2024. [Online]. Available: <https://arxiv.org/abs/2407.01896>
- [248] A. Fan et al., "Large language models for software engineering: Survey and open problems," 2023. [Online]. Available: <https://arxiv.org/abs/2310.03533>
- [249] X. Hou et al., "Large language models for software engineering: A systematic literature review," 2024. [Online]. Available: <https://arxiv.org/abs/2308.10620>
- [250] M. Zakeri-Nasrabadi, S. Parsa, M. Ramezani, C. Roy, and M. Ekhtiarzadeh, "A systematic literature review on source code similarity measurement and clone detection: Techniques, applications, and challenges," 2023. [Online]. Available: <https://arxiv.org/abs/2306.16171>
- [251] K. Wang et al., "Software development life cycle perspective: A survey of benchmarks for code large language models and agents," 2025. [Online]. Available: <https://arxiv.org/abs/2505.05283>
- [252] Hugging Face, 2016. [Online]. Available: <https://huggingface.co>
- [253] S. Chen, P. Puzarla, and B. Ray, "Dynamic benchmarking of reasoning capabilities in code large language models under data contamination," 2025. [Online]. Available: <https://arxiv.org/abs/2503.04149>