

Review Article

Challenges and Paths Towards AI for Software Engineering

Alex Gu¹, Naman Jain², Wen-Ding Li³, Manish Shetty², Yijia Shao⁴, Ziyang Li⁵, Diyi Yang⁴, Kevin Ellis³, Koushik Sen², Armando Solar-Lezama¹

1. CSAIL, Massachusetts Institute of Technology, United States; 2. University of California, Berkeley, United States; 3. Cornell University, United States; 4. Stanford University, United States; 5. University of Pennsylvania, United States

AI for software engineering has made remarkable progress recently, becoming a notable success within generative AI. Despite this, there are still many challenges that need to be addressed before automated software engineering reaches its full potential. It should be possible to reach high levels of automation where humans can focus on the critical decisions of what to build and how to balance difficult tradeoffs while most routine development effort is automated away. Reaching this level of automation will require substantial research and engineering efforts across academia and industry. In this paper, we aim to discuss progress towards this in a threefold manner. First, we provide a structured taxonomy of concrete tasks in AI for software engineering, emphasizing the many other tasks in software engineering beyond code generation and completion. Second, we outline several key bottlenecks that limit current approaches. Finally, we provide an opinionated list of promising research directions toward making progress on these bottlenecks, hoping to inspire future research in this rapidly maturing field.

Naman Jain, Wen-Ding Li, and Manish Shetty equally contributed to this work.

1. Introduction

AI for software engineering has made remarkable progress recently, becoming a notable success within generative AI. Despite this, there are still many challenges that need to be addressed before automated software engineering reaches its full potential. With additional efforts, it should be possible to reach high levels of automation where humans can focus on the critical decisions of what to build and how to balance difficult tradeoffs while most routine development effort is automated away. Reaching this level

of automation, however, will require substantial research and engineering efforts across academia and industry. This paper provides an opinionated view of the tasks, challenges, and promising directions towards achieving this goal.

Many existing surveys overlap with the topics that are discussed in this paper.^[1] and^[2] survey the successes and challenges of AI programming assistants,^[3] survey using LLMs for software testing, and^[4] survey using LLMs in low-resource and domain-specific languages, and^[5] focus on automated program repair, both with and without LLMs. Finally,^[6] is a roadmap for formal mathematical reasoning and has some overlap with our discussion on software verification.

In addition, many papers discuss the current state, challenges, and future of AI for software engineering^{[7][8][9][10][11][12][13][14]}. Our work draws inspiration from them, and we recommend that the reader consult with them for alternative perspectives.

In this paper, our goal is threefold. In Sec. 2, we provide a structured taxonomy of concrete tasks in AI for software engineering. In particular, we emphasize that there are many other tasks in software engineering beyond code generation and code completion, encouraging research in these areas. We provide three measures for categorizing concrete realizations of each task: the scale of the problem, the logical complexity, and the level of human intervention.

Moving forward to Sec. 3, we highlight nine challenges in the field that today's models face, each cross-cutting and applicable to several tasks. In Sec. 4, we posit a set of promising research directions to tackle the challenges above, with Fig. 1 summarizing which directions correspond to each challenge. We hope that through our paper, the reader can appreciate the progress the field has made, understand the shortcomings of today's state-of-the-art models, and take inspiration from our suggested future ideas for tackling these challenges.

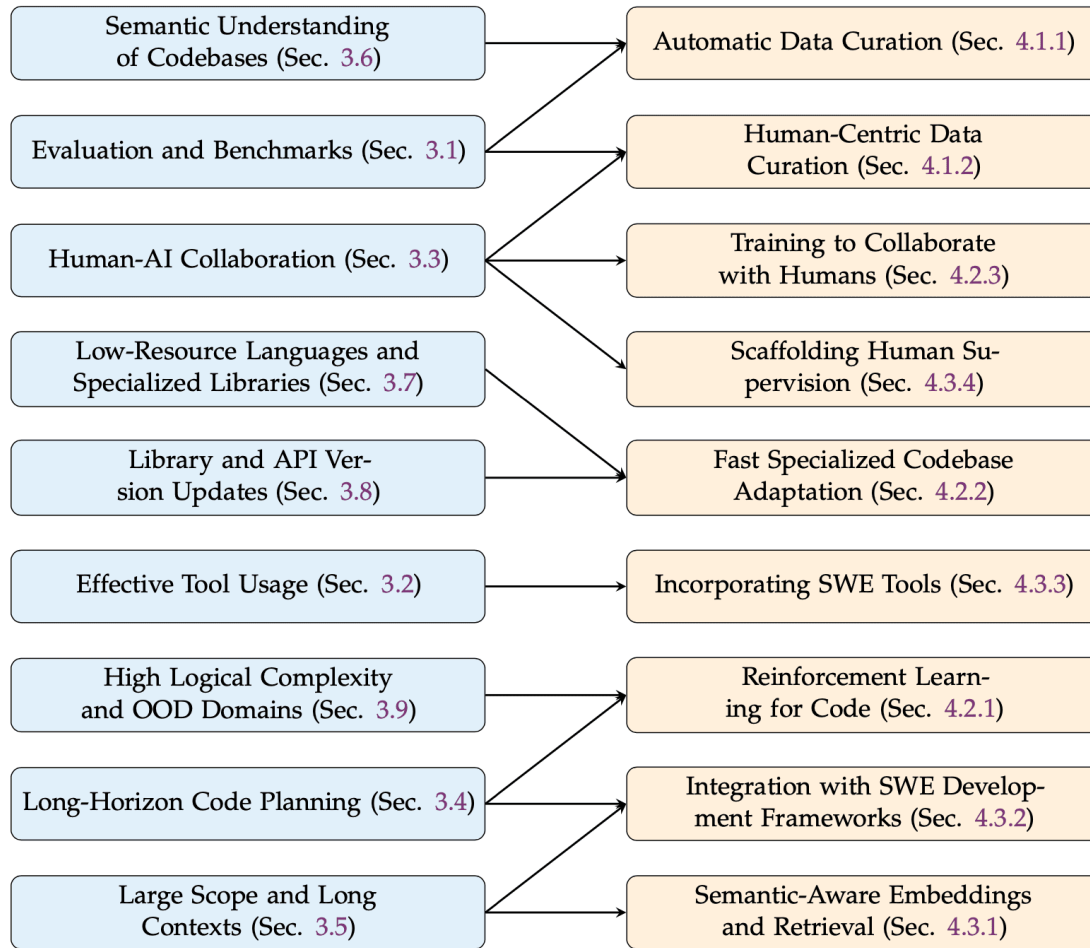


Figure 1. Overview of Challenges (Sec. 3) and Paths Forward (Sec. 4) in AI for Software Engineering

2. Tasks in AI Software Engineering

We first provide a taxonomy of tasks in AI software engineering. To provide a structured way to consider concrete realizations of each task, we define three measures that apply across them: scope, logical complexity, and level of human intervention. To achieve an AI software engineer, we strive for AI to be capable across the board for all three measures.

Scope Measure: We define three levels of scope, the extent of changes that the AI makes to the codebase. *Function-level* scope refers to single, self-contained functions such as in HumanEval^[15] and MBPP^[16]. *Self-contained unit* scope goes beyond singular functions and to larger chunks of code such as entire files and classes, such as FullStackBench^[17] and BigCodeBench^[18]. Finally, *project-level* scope refers to larger codebases such as entire repositories, such as in Commit0^[19] and SWE-Bench^[20].

Logical Complexity Measure: Tasks require a wide range of reasoning abilities when it comes to devising algorithms to solve them. *Low logical complexity* tasks require little to no reasoning, such as writing CRUD (create, read, update, delete) applications or using APIs. *Medium logical complexity* tasks include most LeetCode problems, finding inputs to fuzz simple programs, and reasoning about execution behavior of multithreaded programs. *High logical complexity* tasks require meticulous and challenging levels of algorithmic and logical reasoning, either because the algorithm is complex or because the problem requires clever thinking and insights. This includes difficult competition programming problems, writing large thread-safe concurrent programs, cracking cryptographic ciphers, and solving SMT-like problems. Many popular coding benchmarks are function-level, medium-high logical complexity, such as APPS^[21], CodeContests^[22], and LiveCodeBench^[23].

Level of Human Intervention Measure: AI coding is a collaborative task.^[24] categorize interactions between developers and AI. Each interaction progresses through four phases: the *trigger* for the interaction, the *AI response* describing the system's output, the *developer response* capturing how developers react to the AI response, and the *output* of the interaction, the exact result. They characterize these developer-AI interactions into eleven types, including autocomplete code suggestions, conversational assistance (e.g., asking a question about a codebase), selection-based enhancements (e.g., refactoring a selected chunk of code), comment-guided prompts (e.g., natural language to code), check correctness, and more.

We map these interactions to the autonomy taxonomy outlined in^[25]¹ to define three levels of human intervention. We distill their six levels of autonomy into three levels: low (*No AI and AI as a Tool*), medium (*AI as a Consultant and AI as a Collaborator*), and high (*AI as an Expert and AI as an Agent*). *Low autonomy* is when the human has full control over the task and uses AI to automate simple sub-tasks. This might look like writing a codebase with tests while leaving small function-level snippets for the AI to fill in. *Medium autonomy* is when there is a similar amount of human-AI collaboration, with interactive coordination of goals and tasks. Here, the human and AI might both suggest refactorings, optimizations during the development cycle. *High autonomy* is when AI drives the interaction and tasks, identifying required changes and the changing demands of the user. The AI would defer to the human only when needed or for a final check, write the code and tests autonomously.

Next, with our taxonomy of measures in place, we turn to the set of tasks that are reflective of the tasks and capabilities of a human software engineer. We give a brief description of each task in this section, deferring a more extensive survey to Appendix A.

2.1. Code Generation

Code generation is the task of generating code from a specification. In *code completion*, the specification takes the form of a preexisting code snippet, and the goal is to complete the snippet. The most popular form of code completion is tab completion, where the user can press the tab key to complete a block of code (e.g. GitHub Copilot). Tab completion is often done at line-level or function-level scopes but needs to be fast to provide users with a seamless experience. Another paradigm is *natural language to code*, where the specification is a natural language description with requirements such as the task description, input-output examples, or libraries to use.

Recently, AI-driven IDEs, such as [Cursor Composer](#) and [Codeium's Windsurf Editor](#), have blurred the lines between the two paradigms. With the ultimate goal of decreasing the burden of human programmers, they aim to automatically infer the user's intent from the code context and user behavior (e.g. keystrokes, user edits, file navigation patterns). However, when intent is vague, they allow users to specify desired functionality via chat interfaces. Depending on scope and logical complexity, code generation can vary highly in difficulty. Reliable code generation in large codebases is still a challenge for state-of-the-art AI systems today.

2.2. Code Transformation

2.2.1. Code Refactoring

In code refactoring, the goal is to take a working implementation of a piece of software and rewrite parts of it while maintaining correctness. One challenge with this task is that success extends beyond functional correctness or metrics. The goal is often to improve code maintainability, readability, or extensibility—qualities that can be inherently difficult to quantify and highly context-dependent.

For instance, extracting shared functionality into helper methods presents trade-offs between modularity and cognitive complexity^[26]. While there are no hard rules for when to extract functionality, one heuristic adopted by software engineers is the rule of three (“three strikes and you refactor”)—abstractions should only be used when a piece of code has been duplicated thrice. Because it can often be unclear what level of abstraction refactorings should be done at, completing a refactoring at a high autonomy level is also difficult. These challenges are further compounded by the need to understand implicit trade-offs customized to specific codebases, respect conventions, and reason about the long-term maintenance implications of structural changes. While code refactoring often has a low logical

complexity, it can be laborious in practice due to scope, as seemingly small refactors can propagate across the entire codebase.

Example: React Fiber architecture refactor: React's major refactoring was motivated by performance limitations in the original engine, particularly for complex UIs with animations and dynamic updates. Beyond challenges related to optimized implementation, a major challenge was providing backward compatibility while completely rewriting React's core algorithm. Being an open source tool, this refactor also required educating developers about new concepts without disrupting existing mental models highlighting nuances in real-world software system design.

2.2.2. Code Migration and Translation

An incredibly resource-intensive (time and manual effort) task frequently affecting companies is migrating large amounts of code while preserving all the original functionality and semantics. Such high-value migrations present opportunities for AI-assisted automation to reduce cost and manual effort. Code migration often has a very high scope (many files and systems affected alongside their interdependencies) and high logical complexity (semantic depth of required transformations, constructs in different languages may be different). Current solutions may excel at migrations with high scope but modest logical demands (API migrations, type conversions) but struggle with changes across component boundaries^[27].

A special case of code migration is code translation (transpilation): rewriting code from a source language to a target language. In industry, this task can be motivated by several reasons, such as security and scalability concerns in legacy languages, avoiding the technical debt a project has accumulated over the years, and improving the performance of a codebase. Due to the safety-critical and cross-system nature of many migrations, this task often requires substantial human oversight in practice and cannot be done fully autonomously.

Example: Scala version migration: A recent Scala 2.13 to 3 migration^[28] illustrates these challenges, documenting a year-long effort. Critical issues included the loss of macro annotations, broken type projections, incompatible libraries, and compiler performance degradation—all requiring innovative workarounds and architectural changes. There have been many similar language migrations with analogous problems, famously Python 2 to 3 and Swift 4 to 5.

Example: COBOL: COBOL powers 80% of in-person financial services transactions and 95% of ATM swipes while processing \$3 trillion in commerce a day, with over 220 billion lines of COBOL code in production^[29]. However, there are less and less COBOL programmers, leading to the desire to migrate out of COBOL and into a modern language like Java^{[30][31][32]}. However, because of the large scope and high logical complexity of existing COBOL code, migrating from COBOL to Java would be a monumental undertaking and many companies opt to continue using COBOL. These companies are still forced to migrate to newer versions like COBOL V6, because early versions of COBOL were gradually withdrawn from service. This task still requires skilled COBOL engineers and high precision, as it can often be difficult to understand the business logic of legacy code and introducing bugs can have dangerous implications.

Example: Twitter migration to improve latency: Twitter^a built its initial platform using Ruby on Rails, facilitating rapid development. However, as the user base expanded, performance and scalability issues arose. They migrated key components to Java and Scala, achieving a 3X latency drop. This transition required re-architecting the system to adapt Ruby's dynamic features to the statically typed environments of Java and Scala, exemplifying the complexities of large-scale code translation.

^a <https://www.infoq.com/news/2012/11/twitter-ruby-to-java/>

Example: C to Rust: There has been a push to use translation as a proactive approach to eliminate memory safety vulnerabilities in C-based systems. This has even garnered attention from the US Department of Defense^a, which has long-lived systems that disproportionately depend on C, supporting programs to translate C codebases to Rust (TRACTOR). Recent efforts like Syzygy^[33], C2SaferRust^[34], and AlphaTrans^[35] have shown the potential for hybrid approaches combining LLMs with traditional program analysis techniques. However, some significant challenges remain, including ensuring correctness in large codebases while maintaining desirable attributes such as speed, reduced vulnerabilities, and idiomatcity.

^a <https://www.darpa.mil/news/2024/memory-safety-vulnerabilities>

2.2.3. Code Optimization

Transforming programs to improve performance characteristics while maintaining functional correctness is a critical software task. Optimizing real-world systems poses significant challenges due to the large scope and high logical complexity of the task, as performance bottlenecks must be identified and new algorithms to mitigate them must be proposed. Code optimization often has a large search and solution space with competing objectives like speed, memory efficiency, and readability, for example when optimizing kernel code at the PTX level for GPU-based AI model optimization^{[36][37]}. In many scenarios, high levels of autonomy may not be desirable, as tradeoffs can depend heavily on external factors such as hardware, and the best optimizations may ultimately affect readability.

Example: Google Chrome performance improvements: For over two decades, changes to the Chrome web browser have been an exemplar of optimization affecting real-world code^[38]. Their V8 engine achieved a 20x performance improvement through coordinated optimizations across multiple layers - from implementing concurrent garbage collection that reduced bloat by 100x to developing specialized compilers like TurboFan that improved performance by 5-10%, to enabling background parsing and compilation that reduced compile time by 20%. The demand for cross-layer and low-level code changes (e.g., writing a new JavaScript interpreter) and building tools to measure and test representative performance metrics are key challenges for achieving this sort of real-world impact with LLMs.

2.3. Software Testing and Program Analysis

In the process of software development, there will inevitably be bugs. The difficulty of detecting these bugs can vary depending on their scope and logical complexity. For LLMs, minor typos or correctness bugs (small scope, low logical complexity) are easier to spot^[39] while complex concurrency bugs and security vulnerabilities (large scope, high logical complexity) can be tricky because they can be hidden deep in the call stack, contain subtle logic errors, or be hard to isolate due to the large scope^[40].

2.3.1. Software Testing

Software testing is a practical approach to prevent bugs, both during development and production. There are several popular approaches to software testing, some short-term and others longer-term. *Unit testing* refers to using input-output style tests that exercise the functionality of a piece of code. *Property-based*

testing is based on formal specifications and relies on specifying test cases that ensure that known properties of the code hold. *Mutation testing* modifies a program subtly and ensures that the test suite can detect errors in these mutations. *Fuzzing* refers to executing programs with unexpected inputs and monitoring for exceptions, usually over a more extended time period.

Example: OSS-Fuzz on FreeType: OSS-Fuzz^[41], Google’s automated fuzzing infrastructure, has proven its value by swiftly uncovering security flaws in critical software. For instance, when a recent source change was made to FreeType—a font rendering library deployed on over a billion devices—OSS-Fuzz detected a heap-buffer-overflow within hours:

```
ERROR: AddressSanitizer: heap-buffer-overflow on address 0x61500000ffa READ of size 2 at
0x61500000ffa thread T0 SCARINESS: 24 (2-byte-read-heap-buffer-overflow-far-from-bounds)
#0 0x885e06 in tt_face_vary_cvtsrc/truetype/ttgxvar.c:1556:31
```

The goal of software testing is to design tests that can surface bugs reliably. This is evaluated through metrics such as code coverage—how much of the source code is executed when the test suite is run. An alternate to code coverage is mutation score, where mutants are generated, and the score is defined as the percentage of mutants causing the suite to fail. While practical, software testing faces challenges such as the scalability limits of traditional tools and the difficulty of manually designing tests with good coverage. As LLMs continue to improve at coding, they present a promising avenue for automatically generating high-quality tests.

Example. Fault-based test generation at Meta: Meta’s Automated Compliance Hardening (ACH) system^[42] is a system that generates tests aiming to catch real-world bugs. ACH works in three steps: first, the engineer describes the bugs they are worried about. Second, ACH combines LLMs with mutation testing to generate code with those bugs. Finally, these mutants were used to develop unit tests capturing them. ACH was used to generate tests for Messenger and WhatsApp, where engineers accepted 73% of its tests.

2.3.2. Program Analysis

While testing catches bugs, the most challenging software issues are security vulnerabilities and zero-day exploits, from memory corruption to privilege escalation. This requires a deep program

understanding, that testing/fuzzing often misses. For instance, a *zero-day* is a vulnerability unknown to the software developers that is found by an attacker, and there is no patch available from the vendor. In such cases, the only practical approach is offensive security research, manual source code audits, and root cause analysis of prior vulnerabilities to harden codebases.

Example: Variant Analysis: Project Zero's^[43] investigations at Google reveal that many in-the-wild 0-day exploits aren't entirely new—they're often variants of vulnerabilities that had been patched before. In their analysis of recent 0-day reports, nearly half of the issues were closely related to earlier bugs (such as those affecting Windows win32k and iOS IOMobileFrameBuffer). This finding underscores the importance of performing rigorous root cause and variant analyses. Instead of just fixing a single exploit path, security teams must comprehensively address the underlying bug class, ensuring that alternate exploitation routes are closed off for good—making this task more challenging.

Another example of a valuable but challenging analysis is invariant detection. A *program invariant* is a property of a piece of code that is guaranteed to be true at a specified program point, no matter what the input is. A simple example is that after the line `int x = c * c;` is executed, `x` must be nonnegative. Identifying invariants in a program can be useful when testing, debugging, and modifying code. This task can be challenging because it requires reasoning abstractly about code execution across many different potential inputs and execution paths to determine what properties must hold for all possible inputs.

2.3.3. Program Repair

Bug localization is a significant challenge in program repair, as pinpointing the exact site of a bug can be challenging, especially in large codebases. Issues like out-of-memory accesses often manifest themselves further downstream, making it difficult to identify the root cause. Once the bug is localized, the next step is to repair the bug. LLMs can be an ideal tool for this because they have seen a wide variety of bugs during training. Function-level, low-logical complexity bugs can often be easily fixed by feeding back error information to the model. It can be trickier to perform repair in larger scopes (e.g. repositories) where the code has higher logical complexity. This can often require several steps, including designing and implementing new algorithms or making complex refactorings across multiple files.

2.4. Software Maintenance

2.4.1. Code Documentation and Summarization

To ensure maintainability, readability, and ease of collaboration, code must be well documented. Good documentation needs to be written cleanly and crisply, describing what the function does and how the function works. It must also anticipate and address any misunderstandings that a programmer might have, such as potential side effects or special cases. Humans often see documentation as a chore and neglect it, leading to code and documentation frequently being out of sync. This has led to the concept of “self-documenting code”, code that clearly conveys its purpose. As documentation is generally a task that has a low logical complexity and does not require too much human intervention, LLMs can help ensure that documentation is a continuously updated artifact in sync with the code.

2.4.2. Pull Request (PR) Review

Reviewing pull requests is an integral aspect of the software development cycle. While the most essential requirement for PRs is that a new feature is implemented correctly, other important considerations include checking whether the repository’s style conventions are satisfied, ensuring that the PR does not introduce any new bugs, verifying that program invariants and guarantees still hold, and inspecting whether tests are robust. Generally, reviewing PRs is a task requiring low logical complexity and can be automated relatively easily.

2.4.3. Code Understanding, Navigation, and Question Answering

When encountering a codebase for the first time, developers often find it challenging to understand and develop a good mental model of the code. This can be due to many reasons: too many wrapper functions, excessive error-handling boilerplate, deep call stacks, or poor code cleanliness. One important challenge in code understanding is code navigation: finding where relevant functionality is implemented. Doing this well requires understanding the high-level layout of where every functionality lies in the codebase and the low-level understanding of which helper functions are used to implement each functionality.

Another challenge is code question answering: answering complex questions about a codebase, which requires sophisticated code understanding and reasoning abilities. Models should not hallucinate or give incorrect information that skews a developer’s mental model of the code. Beyond other tasks mentioned in this section, developers might commonly ask questions related to data flow (when and where data

structures get mutated), code functionality (whether there are any side effects), performance characteristics (determining the runtime and memory complexity of a function), or error handling (whether certain corner cases are handled).

2.5. Scaffolding and Meta-Code

For a software system to work, the core logic must be written well, but that is not enough. Many infrastructural aspects must be in place to support the software. We group these into two main categories: scaffolding and meta-code. We define scaffolding as a task outside of the code that must be done to get the software running properly. Examples of scaffolding include setting up Google authentication, subscribing to APIs, managing file storage, and generating API tokens. In contrast, we define meta-code to be code that is important to make the system work but does not actually participate in the execution of its main logic. Examples of meta-code include test harnesses, configuration files, CI/CD code, Makefiles, Dockerfiles, sandbox databases, and preprocessors. Scaffolding and meta-code often are small in scope and have low logical complexity but can require a lot of domain-specific knowledge about the application, requiring human intervention.

Example. Configuration validation: Ciri^[44] is a tool that uses LLMs for configuration validation on open-source projects including Django, PostgreSQL, and Redis. They find that while Ciri excels at detecting misconfigurations of syntax and range violations, it struggles to detect dependency and version violations and is limited to a narrow range of misconfigurations. They also find that LLMs are biased towards more popular configuration parameters, which may lead to hallucinations in out-of-domain scenarios.

Infrastructure-as-code and Security. A particularly challenging case is generating *Infrastructure-as-code* such as Terraform, where you specify the type of infrastructure specifications (such as AWS EC2 instances, Kubernetes clusters, S3 buckets, VPC buckets) as code and execute it to create the infrastructure. When generating such code, LLMs struggle with security configurations due to the complex interplay between service-level permissions (e.g., AWS resource access), resource-level permissions (e.g., specific allowed actions), and provider-specific security primitives like IAM roles, security groups, and network access controls.

Example. Distinguishing permission levels in cluster setup: ^[45] show that on a task of bringing up a cluster, models fail to distinguish between ECS (Amazon Elastic Container Service) Task Execution Roles (for container operations) and Task Roles (for application-level permissions). This resulted in overly permissive policies where a single role was granted both image pull permissions and DynamoDB table access, violating principles of least privilege.

2.6. Formal Verification

The task of formal verification involves generating checkable, mechanized proofs that can guarantee that a piece of code works as intended. There are two major classes of formal verification: full functional verification (FFV) and property verification (PV). In FFV, the goal is to design a complete and precise formal specification that captures the desired behavior of the implementation, such as fully verified data structures (mutable lists, trees, graphs, hash tables)^[46]. The main challenge in full functional verification is in correctly writing the specification so that all desired properties are specified. FFV generally has a high scope and medium logical complexity, as the properties to verify are often straightforward to write once the correct abstractions are found.

While FFV provides a complete set of guarantees, it is usually sufficient to opt for PV, where a few key properties of a system are proven correct. Examples include: ensuring that two threads do not simultaneously enter a critical section of a program, verifying that a complex program will always terminate, proving the absence of security vulnerabilities like buffer overflows, and guaranteeing memory safety. One challenge that makes PV difficult to use in practice is the issue of false positives, where functionally correct code often does not pass property checks. A prime example is Rust: while the powerful type system enforces many desired guarantees, code with correct semantics often does not pass type checks. Another challenge is that many standalone tools for PV are often semantics-dependent, which can make them hard to maintain as language semantics change.

Example. Costly disasters: Formal verification of software is important in mission-critical applications such as aircraft software, as software bugs may lead to costly disasters. In the maiden flight of the Ariane 5 rocket, a floating-point conversion error caused it to explode forty seconds after liftoff. Another case is with the computer-controlled radiation therapy machine Therac-25, where concurrency bugs led to six people being massively overdosed, leading to serious injury and deaths.

Example. Verified Compiler: CompCert ^[47] is a formally verified optimizing C compiler that supports a restricted subset of C including most of the ISO C 99 language. CompCert has been formally verified using the Coq proof assistant ^[48], eliminating the potential for compiler bugs.

While formal verification tools have begun to see adoption in industry, they have not yet become mainstream because of these challenges. Code LLMs could greatly ease this burden and make it more feasible to verify code at larger scales, especially verifying properties requiring lower logical complexity.

Example. Property Verification: Coverity: Coverity is a static analysis tool meant to find generic errors (memory corruption, data races) and system-specific violations (e.g. function-ordering constraints). In their report ^[49], they highlight two issues mentioned earlier: churn and false positives. The first issue, churn, deals with ensuring that the tool produces the same result both when the code base is modified and across different versions of the tool, making upgrades “a constant headache”. The second issue is that when the false positive rate is more than 30%, users ignore the tool and real bugs get lost among these false positives.

3. Challenges

While the field of AI for code has made fruitful progress, cutting-edge AI still struggles with SWE tasks, especially at larger scopes and higher levels of logical complexity. Next, we discuss ten key challenges in AI for code. Each challenge spans multiple tasks, and progress on any can lead to improvements on many tasks at once.

3.1. Evaluation and Benchmarks

Today’s code LLM evaluations focus on a narrow set of tasks, suffer from potential contamination, and do not reliably measure real-world software engineering abilities.

Potential solutions: 4.1

Our taxonomy of tasks and measures highlights some of the shortcomings of today’s evaluations and benchmarks. For example, the majority of today’s coding evaluations have no level of human intervention, with a few, such as Copilot-Arena^[50], having low to medium autonomy. HumanEval, MBPP,

APPS, CodeContests, and LiveCodeBench are all at function-level scope, with low to medium-high logical complexity. Commit0^[19], SWE-Bench^[20], TestGenEval^[51], RefactorBench^[52], SWE-Lancer^[53] are at project-level scope with low to medium logical complexity.

Task Diversity and Capability Isolation: Current coding evaluations primarily focus on the code generation task, while most of the tasks discussed in Section 2 are either not studied such as Code QA or only studied in limited scopes like EvalPerf^[54], vulnerability detection^[55], formal verification^[56]. As more agent-based approaches are introduced for software engineering (e.g. pairing a code generation model with a debugging model), these engineering-related capabilities beyond just code generation will be important towards designing a maximally performant system. Solely relying on end-to-end coding evaluations that focus on the overall correctness of a codebase makes it difficult to precisely measure progress and learn from the failure modes on individual tasks.

Contamination: Data contamination is a serious issue that, if not taken into account, can affect the soundness of various conclusions drawn from a set of benchmark results. In coding, the performance of LLMs on competitive programming^{[57][23]} and SWE-Bench^[58] tasks has been shown to degrade over time, indicating the possibility of older problems being contaminated due to public exposure on the internet. For simpler function-level HumanEval style problems,^[59] suggest three potential causes of contamination: direct data leakage (benchmarks are on GitHub), synthetic data leakage (there are only a limited number of interview problems), and overfitting to test sets (benchmark hacking). In addition, for code, contamination can be hard to detect, as semantically equivalent code that is syntactically distinct could be thought of as contamination^[60]. A recent benchmark, the Konwinski Prize², is a promising way to fairly evaluate SoTA LLM models by only using new GitHub issues.

Construct Validity: Construct validity refers to how closely a measurement reflects the underlying concept. Given the implications of rapid performance improvement in the AI for the code domain, it is essential to have high-construct validity benchmarks evaluating how well programming agents can perform. While benchmarks like SWE-Bench come close, user experiences do not currently match rapid performance gains obtained from them. This is partially because many desiderata in software engineering cannot be described cleanly via automated unit testing. Things like multi-turn code generation, designing an UI, and writing clean and idiomatic code are all difficult to quantitatively measure with precision. Designing reliable proxy metrics for these desired goals remains a challenge.

3.2. Effective Tool Usage

While software engineers use a wide suite of programming tools when programming, most of today's AI coding systems do not invoke tools. AI needs to be able to select which tool to use, decide how to use it, and interpret the outputs in order to continue making progress on the task.

Potential solutions: 4.3.3

Software engineering has witnessed the development of various open and proprietary tooling support for programming, debugging, analysis, and code management over time. For example, program analysis tools provide static and dynamic assurances on code correctness. Print statements and debuggers are used for dynamically analyzing and debugging programs at a fine-grained level. Beyond programming, such tools are richly integrated into the entire software development lifecycle, e.g., code navigation or search, reviewing code, CI testing.

There have been efforts combining LLMs with tools such as calculators and search engines^{[61][62]}. However, effective integration of LLMs with software engineering tools is a more challenging problem. Several early works have incorporated such tool feedback in code generation in an automated fashion, for example, linter or execution feedback in^{[63][64][65]}. However, these works do not actively interact with tools. More recently, programming agents have started incorporating tool use within their workflows termed as Agent-Computer-Interface^[66]. These tools range from aiding in general search (grep), providing code editor for making changes^{[67][68]}, language server for static analysis^[69], dependency analyzer^[70], terminal access for bash commands including code execution^[66], debugger^[71].

Dynamic and Effective Tool Usage: While many efforts combine LLMs and agents with tools, they do not achieve fully dynamic and effective software engineering tool usage. This involves an AI system seamlessly and proactively integrating appropriate tools depending on the task at hand. There are a few challenges towards achieving this goal. First, the AI system must identify which tools could potentially be useful for the task at hand. Second, the system then needs to decide when to invoke the tool. A complex debugging task might require the use of pdb or gdb to track intermediate program states, while looking at input-output pairs may be sufficient for simple debugging tasks. Third, the agent then must figure out how to invoke the tool. If the agent knows that a certain function in a program has an error, it may wish to walk through only that function instead of the entire code from start to finish. Finally, the agent needs

to incorporate the output provided by the tool in order to inform its next steps, e.g. edit the code if a bug was uncovered or run the tool again otherwise.

Example: Performance Instrumentation: A common way to instrument software systems is known as *compiler-inserted program instrumentation*. CSI^[72] is a tool that inserts instrumentation hooks to track objects such as memory loads/stores, function entry/exits, and basic blocks. CSI contains tools like code coverage tools, a memory-operations counter, a performance profiler, and a call-graph generator. To use the tool, the user must follow the API in order to write hooks so the correct aspects can be profiled. Tools like CSI are very valuable when trying to improve the performance of a piece of code, but are not trivial to use. In order for an LLM agent to use CSI effectively, it must first familiarize itself with the CSI API. Then, it needs to know exactly which aspects of the code to instrument, such as placing hooks before and after a function suspected to be a bottleneck. Finally, the agent needs to learn how to use the output of the tool to inform its approach to the task, such as deciding whether a block of code can be further optimized after seeing its performance profile.

3.3. Human-AI Collaboration

Human-AI collaboration is still far from seamless. First, specifications written by humans are often vague and leave out many details, leading LLMs to produce code misaligned with humans. There is also very little controllability with coding LLMs, and today's human-AI collaboration interfaces are still limited.

Potential solutions: [4.1.2](#), [4.2.3](#), [4.3.4](#)

While AI coding systems are increasingly more powerful, the majority of them are still at a low to medium autonomy level, serving as engineer assistance rather than achieving high or full automation. We identify a few key challenges of today's AI coding systems that prevent these systems from working with humans effectively at higher levels of autonomy.

Vague Specifications and User Misalignment: When using code LLMs or coding agents, we typically prompt them with a natural language specification. This can include a natural language description of the desired code, input-output examples, relevant code snippets, and other functional requirements. However, there is a gap in the level of abstraction between English and code, leading to incomplete or ambiguous specifications. This issue becomes more pronounced in longer programs, where the number of ambiguous decision points increases, and choices traditionally made by humans are instead implicitly

embedded in the LLM's generated code. Consequently, users often experience misalignment due to vague specifications. While many code LLMs support multi-turn interactions, it remains inherently challenging for users to articulate their thought processes into follow-up natural language instructions.

Specifications beyond text: While today's specifications predominantly rely on text, there are many domains for which pure text is insufficient as a specification. In domains like robotics, virtual reality, embedded devices, and user interfaces, specifications often need to be multi-modal (e.g. showing the model a picture of an UI to create) and world-interfacing (e.g. providing simulation code describing a robot will interact with its environment).

Inherent trade-offs in software development: Designing large software systems always surfaces trade-offs between various desiderata such as readability, scalability, performance, maintainability, reliability, security, etc. These trade-offs are often context-dependent. A long-term and rapidly moving project may be willing to trade off some performance to have simplicity and readability. Performance-critical applications may completely sacrifice readability to eke out every millisecond of speed (such as using bit-twiddling hacks). Finding the sweet spot among these trade-offs can often involve extensive prototyping and benchmarking to understand the performance characteristics of different approaches. However, user specifications in the initial prompt rarely include details about these trade-offs, nor do models often take them into account.

Implicit constraints: Aside from functional/semantic correctness, there are also often implicit constraints in writing code. For example, many companies such as Jane Street and Google have style guides, and many GitHub repositories explicitly outline style elements that new code ought to follow. ^[73] find that GitHub pull requests that are more consistent with the style of the existing code get merged faster. Additionally, corporations may enforce codes of conduct or compliance requirements at the code level. Furthermore, codebases follow common programming patterns or system design patterns that are implicitly specified by the way the current code is written. However, when using code LLMs, these constraints are often inferred incorrectly^[74].

Example: Serializer-Deserializer pattern for objects: Consider the issue [astropy-#14181](#) from the astropy Python library. The issue requests support for a new input file format (reStructuredText) to load astronomical data into the codebase more flexibly. While the issue does not mention it explicitly, as per common practices, developers implement read (deserializer) and write (serializer) operations when implementing support for a new file format. This ensures data can flow bidirectionally between the file format and the application's internal data structures. However, models evaluated on this issue, as part of the SWEBench benchmark, only implemented the read method.

Lack of Controllability: When using AI coding systems, programmers often seek specific functionality, yet they lack reliable ways to steer LLMs toward generating precisely the desired code. Instead, they typically rely on a trial-and-error approach, repeatedly sampling outputs or providing feedback until the AI produces an acceptable solution. Consequently, significant human effort is still required to review and modify the code to ensure it meets the intended functionality^[75].

A way to improve controllability is for AI coding systems to recognize when human input is needed and communicate effectively—yet this remains the top-reported challenge in human-agent collaboration^[76]. LLMs rarely defer to humans for clarification, while developers often ask questions to clarify the description of a task provided by their peers. For example, when a product manager refines a requirements document, developers who are unclear about the scope or specifications ask questions and leave comments, which the manager resolves iteratively to disambiguate requirements^[77]. Based on its knowledge of existing software, AI should be able to incorporate implicit priors from a specification while keeping the user in the loop. For instance, when designing an academic website, certain expectations—such as including a list of publications and contact information—are implicit. However, whether to include a person's GPA would require explicit clarification.

Restricted Human-AI Interface: Existing interfaces for code LLMs primarily manifest as intelligence features embedded within integrated development environments (IDEs).^[24] establishes a taxonomy of developer-AI tool interactions, emphasizing low-level support mechanisms such as auto-complete suggestions, selection-based enhancements, and conversational assistance within the codebase context. While this taxonomy comprehensively covers existing AI coding systems that function primarily as engineering assistants, its applicability becomes questionable as these systems advance toward higher levels of automation. For instance, the ubiquitous “Tab” interaction paradigm prevalent in intelligent

IDEs may prove inadequate when AI systems transition from completing developer-scaffolded functions to authoring the majority of the codebase autonomously.

Current interfaces for coding agents, such as Devin, typically stream raw actions without adequate context or explanation. Given that these agents can execute numerous actions rapidly, developers face significant challenges in effectively monitoring the process, implementing timely interventions, or reasserting control when necessary. This lack of transparency can also undermine trust in AI-generated code^[78]. While human-AI interface design has received extensive attention in autonomous vehicle research^{[79][80]}, similar consideration for AI coding systems remains notably absent.

3.4. Long-Horizon Code Planning

Large software engineering projects often require long-term planning about the design and structure of the code. LLMs struggle at two key aspects of this: designing good, lasting abstractions and respecting modularity and code quality principles.

Potential solutions: [4.2.1](#), [4.3.2](#)

When working on large projects, engineers and tech leads often make complex decisions about how to design and structure the code to best support the various functionalities that will eventually be needed. To build a long-lasting software system, an engineer must know the potential paths that the system's evolution might take. This requires domain expertise and experience in how different code structures require different forms of extension. We believe that today's language models are unable to perform this level of sophisticated planning.

Designing Good Abstractions: One instance of long-horizon code planning is choosing the right abstractions from the outset. An API designed with good abstractions will allow new features to be implemented seamlessly with minimal user overhead, while an API designed with poor abstractions may lead to excessive code duplication, refactoring, or even debugging. We discuss two examples of this, library learning and data representation.

Library learning: Designing APIs and libraries are designed with useful abstractions often leads to more code reuse and more intuitive interfaces. The challenge of library learning is to derive a library of useful abstractions from a corpus of programs by abstracting out common reusable features^{[81][82][83]}. While the

traditional library learning literature has focused primarily on code reuse, a truly effective library must also prioritize ease of use and maintainability, as well as be robust and adaptable to future extensions.

Data representation: The choice between data structures leads to a variety of trade-offs when it comes to performance aspects such as memory usage and processing speed. For example, database engineers need to decide between various data models, storage formats, and indexing methods to balance performance.

Example: Database Design for Web Applications: Database engineers strive to design their databases in a way that minimizes memory usage and maximizes query performance (speed). To achieve this goal, the databases community has spent considerable efforts optimizing both the high-level data representation and the underlying data structures^{[84][85]}. Consider the task of designing a database schema for a restaurant owner to manage their business: keeping track of customers, managing a rewards program, maintaining the restaurant's inventory of ingredients, etc. One important design decision to make is deciding on a schema: while having a reservation and customer table is fairly straightforward, should we include a table for customer reviews or simply add rating and review fields in the customer table? Another important design decision is choosing which database indexes to include. While choosing the correct indexes can speed up queries significantly, indexes cost additional memory and must be kept updated. Making decisions like these requires knowledge of the application, context, and the effects of each option.

Modularity and Code Quality: LLMs are trained and optimized primarily for code correctness with insufficient focus on other aspects of code like quality and maintainability. This is further exacerbated with large scale reinforcement learning being performed using test cases which can lead to unintended consequences regarding code quality, as correct but poor quality code is still often given a high reward. Empirically, it has been observed that LLM written solutions are often more complex than human-written counterparts. For example,^[86] identified that LLMs prefer to repeat existing code instead of making use of abstractions in the existing code. One aspect of code quality is modularity, ensuring that code does not get duplicated too often. Here,^[87] identified that library or tool reuse is non-trivial for LLMs in coding and formal math domains.

3.5. Large Scope and Long Contexts

Coding is a domain where required context lengths can be very long, as codebases can consist of millions of lines of code, posing challenges to today's models. In addition, today's retrieval-based methods are still limited: models often retrieve incorrect information and can still struggle with leveraging retrievals effectively due to the difficulty of code reuse.

Potential solutions: [4.3.1](#), [4.3.2](#)

Large Scopes: At the repository level, the tasks in Sec. 2 become significantly more difficult and require many steps. In code generation, user alignment can be an issue because there are many decision points and tradeoffs that can compound. In code refactoring, modifications will touch multiple parts of the codebase, and it can be tricky to keep the repository consistent. In code debugging, functions can be large and bugs can be nested deeply within stacks of function calls. In code navigation, because there are so many functions interacting in various ways, it can be difficult to know where each piece of functionality is implemented and how the code is pieced together.

Another issue with large scopes is large context lengths. Software engineering often requires dealing with very large codebases—for example, Google has repositories with over a billion lines of code^[88]. As this is far too large for modern-day LLMs, choosing the correct context to include when using LLMs is important.

Example: Debugging Cloud Applications: Organizations often rely on monitoring and observability tools to track the performance of their applications. One such tool is Datadog, an observability service for cloud applications that can monitor infrastructure, detect security anomalies, and track database performance. For larger applications with more moving parts, these logs can consist of thousands of lines of JSON payloads. For humans, sifting through these logs is usually a matter of searching for certain keywords that they know will appear in the logs. However, LLMs often have a hard time interpreting large amounts of logs like these.

Limits of Retrieval-Augmented Generation (RAG): Retrieval-based algorithms have been the predominant way to deal with long-context coding issues. First, the retriever retrieves relevant functions. Then, the generator leverages the retrieval to improve generation. While RAG has proven effective in

many NLP tasks such as question answering^{[89][90]}, the code domain provides new challenges for these methods.

Retrieval: In most NLP tasks, the retrieval step can be done relatively well because keywords that are in the query are often keywords that need to be retrieved. Unlike answering NL questions, writing code often requires drawing inspiration from code snippets that may be completely different syntactically. This can include programs with similar semantics, algorithms, or API calls, all of which potentially have very little in common when it comes to syntax. For example, the implementation of Dijkstra’s algorithm in a GPS navigation application can guide the implementation of a shortest-path algorithm in a social media application. Because retrievers often rely on syntactic matching, these relevant programs can be hard to retrieve^{[91][92]}.

When deciding what to retrieve, it is also necessary to have a sufficient awareness of other parts of the codebase so that you know which building blocks are necessary to construct the new function.

This can make the retrieval task relatively tricky, as shown by failure modes on two benchmarks, CodeRAGBench^[93] and BRIGHT^[94].

Example: Failure Case of Finding Relevant Files When Resolving Issues: BM25, despite its widespread use in code retrieval, demonstrates limitations in scenarios that involve large and complex codebases. For instance, in [chartjs__Chart.js-7951](#) from SWE-bench Multimodal^[95], BM25 retrieval using the issue description returns suboptimal results. The top-3 retrieved files from `src/` are `src/scales/scale.radialLinear.js`, `src/scales/scale.linearbase.js`, and `src/helpers/helpers.canvas.js`. However, the critical modifications required to resolve the issue should occur in `src/elements/element.bar.js` and `src/controllers/controller.bar.js`. This retrieval failure impedes the effectiveness of coding agents, many of which are augmented with code retrieval systems. When agents focus their attention on irrelevant files, their ability to resolve the issue successfully becomes substantially compromised.

Generation: In NLP tasks, the generation step often is a straightforward application of the retrieved information. However, in code, writing a new function requires more than copy and paste. This is closely tied to the problem of code reuse: piecing together relevant snippets of code in a precise and productive way to fit the current context. Depending on what is retrieved, each piece of retrieved content provides different information. This can include information about the language’s syntax, documentation about

the API, clues about the algorithm to be written, or examples of similar functionality being written. [96] find that even when the oracle context is retrieved, LLMs tend to misuse it, highlighting a lack of semantic understanding, which we discuss in the next section.

Example: Bad Generation Despite Identifying the Correct Context: [96] highlights a failure case where a code LLM fails to complete a Python test case correctly, even though it has the correct context. The function name from the context, `test_case_convert_camel_to_snake`, suggests that the function being completed is a test case for `convert_camel_to_snake`. With the given context, the model generates the function as `convert_camel_to_snake`, which however does not match the larger codebase as other pieces of code expect this function name to be `camel_to_snake`. While this issue can partly be attributed to incomplete retrieval of relevant information, it also presents a challenge for code LLMs, as they must recognize such inconsistencies—especially when the immediate context is correctly provided—thereby avoiding high-confidence errors.

3.6. Semantic Understanding of Codebases

Being able to effectively write code relies on having a strong semantic understanding (somewhat like a world model) of the entire codebase: structurally seeing how various parts of the code go together, knowing what is implemented where, understanding how the algorithms work, and keeping track of program invariants at certain program points. LLMs struggle with this global semantic understanding

Potential solutions: 4.1

A global and holistic semantic understanding of a codebase is important for performing almost all code-related tasks. For example, let's say an engineer is asked to improve the runtime performance of a query. To do so, they must first understand the codebase's structure well enough to know where all the pieces of the algorithm are implemented. Then, they need to understand the algorithm and implementation in detail. This includes both the high-level algorithm (including its time complexity) and the low-level implementation details to identify both algorithmic and implementation bottlenecks. Finally, after coming up with a solution, they must then return to their understanding of the global code structure so that they can integrate their new algorithm without introducing new bugs.

LLMs struggle at semantic understanding of codebases for several reasons. First, the way that code is pieced together can be relatively intricate, and understanding all these complex relationships can be

difficult. Second, code can often have units with high logical complexity that contain custom algorithms that may never have appeared anywhere in the training data. Finally, because a disproportionately large number of LLM training tokens are spent on code generation rather than other coding tasks, models may lack a holistic awareness and world model of code.

One desiderata is that models can generalize knowledge across various coding tasks^[97]. However, this may not be straightforward as just training on more tasks:^[98] found that coding models fine-tuned on additional natural language/code pairs saw significant improvements on code generation but did not transfer to improving code understanding and execution reasoning. While there have been successful efforts to augment code LLM training with execution information to improve general coding capabilities^{[99][100]}, imbuing code LLMs with a general and holistic understanding of code remains an important challenge today.

3.7. Low-Resource Languages and Specialized Libraries

Models struggle with low-resource languages and codebases with specialized libraries. Because of the limited exposure they have in these contexts, models fail in a variety of ways including generating syntactically incorrect code, misunderstanding the semantics of certain constructs, and using libraries improperly.

Potential solutions: 4.2.2

As we adapt code LLMs to individual codebases, generating correct code in out of distribution (OOD) scenarios becomes crucial. Much of software development in business contexts revolves around proprietary codebases, which is a distribution shift from the open-source code that dominates LLM training data^[101]. These OOD scenarios include domain-specific languages (DSLs), custom internal libraries, low-resource APIs, and company-specific coding styles and conventions.

Syntactic Failures: Models have been shown to hallucinate constructs from higher resource languages when working in low-resource languages.^[102] remark that "contemporary LLMs fail to follow Hazel's syntax and semantics, often borrowing syntactic forms and library functions from [higher-resource languages like] OCaml and Elm".

Example. Syntax error in Triton: In Listing 3.7, we show an attempt from Gemma-3 27B to write a dot-product kernel in a low-resource GPU programming language called Triton (part of docstring omitted for brevity). Gemma uses indexing notation such as `a[index]`, which is not a valid Triton construct. Models like o1 and o3, however, do not make this mistake.

```
@triton.jit
def dot_product_indexed_kernel(a, b, indexes, out, N):
    """
    Computes the dot product of two vectors a and b using an index vector.
    ...
    """
    block_size = 64 # Adjust block size for optimal performance
    grid_size = (N + block_size - 1) // block_size

    block_id = tl.program_id(0)
    start = block_id * block_size
    end = min(start + block_size, N)

    accumulator = tl.zeros(1, dtype=tl.float32)

    for i in range(start, end):
        index = indexes[i]
        accumulator += a[index] * b[index]

    tl.store(out, 0, accumulator)
```

Poor Semantic Understanding: In low-resource languages, models have less exposure to the various language constructs. Therefore, they have a weaker semantic understanding of the language. Many studies reveal that code LLMs perform poorly when asked to write code in low-resource languages. Due to the lack of training data in these OOD domains, models may struggle to write common primitives or piece together functionality coherently. On HumanEval, Qwen 2.5 Coder Instruct (32B)^[103] has an accuracy of 83% in Python but only 27% in D.³

Library Usage Failures: In OOD scenarios, LLMs lack awareness of the libraries and functions available for use. In new codebases using custom libraries, many functions appear only a few times, providing limited training data for AI models to learn their usage. This scarcity can lead to overfitting, where models fail to recognize an effective use-case of these functions. Models also frequently hallucinate non-existent functions based on patterns that it infers.

Example. Hallucinating theorems in Lean: In Listing 1, the LLM makes up a theorem named `infinitelyManyPrimes_arithmeticSequence`. While the model understands that it should use Dirichlet’s theorem, the actual theorem is named `forall.exists_prime_gt_and_eq_mod`. This theorem occurs very few times in all of the currently available Lean code on GitHub, so it is likely that LLMs did not train on it many times.

```
theorem infinitely_many_primes_of_form_6k_plus_1 :
  ∀ n : ℕ, ∃ p : ℕ, p > n ∧ Prime p ∧ ∃ k : ℕ, p = 6 * k + 1 := by
  intro n
  -- Use Dirichlet's theorem for arithmetic progressions
  have h := Nat.infinitelyManyPrimes_arithmeticSequence 6 1
```

Listing 1: LLM-generated Lean code snippet that incorrectly uses indexing notation

3.8. Library and API Version Updates

Software engineering has a unique property that repositories and libraries are constantly changing. Code LLMs have trouble adapting to these rapid changes, often struggling to use the correct version of libraries and ignoring new paradigms and features.

Potential solutions: [4.2.2](#)

Continual learning, the idea of training an AI system to take in new information continually, has been a long-standing challenge in AI and NLP^{[104][105]}. In software engineering, codebases are continuously changing as new features are supported and awkward design patterns are reworked. While backwards compatibility is often prioritized in software design, it inevitably becomes broken as codebases evolve further. Therefore, programming libraries have version releases, each release supporting and deprecating features in the last version.

There have been a few works exposing this issue. For example, `CodeUpdateArena`^[106] and `GitChameleon`^[107] are two benchmarks exploring the ability of LLMs to write version-specific code, examining this issue at the function and file level. They find that language models struggle to adapt to these changes even with this limited scope. In theorem proving (Lean),^[108] try to mitigate this by developing a lifelong learning framework that continuously learns and uses new theorems. In real-world engineering, the challenge of library and API versioning generally spans across an entire repository, as everything must be kept consistent. To our knowledge, there are no techniques that successfully deal with this challenge at such a large scale. This problem is difficult for a few reasons, which we discuss below.

Version Identification: In order to successfully deal with version changes, a LLM must first identify which version of each library is being used in a codebase. This may often be quite difficult, because versioning information can be hidden deeply within a codebase. Sometimes, it can be found in comments or configuration files, but in the worst case, it must be inferred from the library calls being used. To make things worse, some code may be compatible across multiple versions, while other code will cause errors only in specific versions. Therefore, the model will often require a deep understanding of both the codebase and the nuances between different versions in order to infer the version at hand.

Example: Debugging Frontend Code: Frontend framework usually has more frequent versions update, making it hard for code LLMs to work with. For example, when helping a user debug the “NextRouter was not mounted” issue, Claude 3.7 tries various solutions without recognizing that the core problem requires importing useRouter from 'next/navigation' instead of 'next/router', a crucial distinction since the user’s codebase leverages App Router in Next.js 13.

Version Adaptation: Many fast-changing libraries are not backward compatible as older features become deprecated. It can be difficult for LLMs to implicitly keep track of which constructs and patterns are associated with each version. Therefore, consistently using constructs from the right version can be difficult. As we will see in the examples below, LLMs often write code that mixes and matches API constructs from different versions of the same library.

Example. Typing Hints: While Python 3.5 required importing types from the typing module, Python 3.9’s PEP 585 enabled direct use of built-in types for generics (e.g., `list[int]` vs `typing.List[int]`). However, language models tend to default to the older typing module syntax.

Continuous Adaptation to Paradigms, Features, and APIs: New styles, patterns, and paradigms are often introduced to replace older, more cumbersome ways to write code. For example, React came out with its *Hooks* paradigm in version 16.8 (2019). Over the next few years, developers transitioned from the old class components paradigm to using hooks, as hooks made code cleaner and more maintainable. Only in early 2023, with the launch of react.dev, were Hooks the default paradigm in the documentation. For language

models, incorporating these features can take a long time, because code in these new paradigms are initially completely absent in the training data and inherently in the low-resource regime. In^[109], the authors find that LLMs fail to utilize security features in compiler and toolkit updates (such as in Java 17), still relying on legacy methods such as insecure random number generation. While it is possible to use retrieved examples and documentation in order to coerce language models to write code using new and updated features, we should strive to create AI coding assistants that can quickly internalize new changes and be able to naturally incorporate new features and paradigms, even without an abundance of training data. For each task, the language model should be able to reason about the best way to write the code, independently of the number of occurrences seen in the training data.

Example. Lean 3 vs. Lean 4: Lean^[110] is a programming language that allows users to write formal proofs of mathematical theorems. In 2017, using Lean 3, enthusiasts implemented a library for mathematics called mathlib, with over half a million lines of code. Because Lean 3 had many shortcomings, Lean 4^[111] was initiated at the beginning of 2021 to address many of these issues. There was a massive undertaking to port all of the mathlib code over to Lean 4, and only in September 2023 was there a stable release of Lean 4, the version of Lean that is predominantly used today. The two versions are generally incompatible. We hypothesize that, due to the recency of Lean 4, most language models have been trained on much more Lean 3 code compared to Lean 4 code. When asking models to generate code in Lean 4, it sometimes generates code with Lean 3 coding conventions. Other times, it uses theorems and lemmas from Lean 3 that are deprecated in Lean 4. In Listing 3.8, we show an example of prompting o3-mini with a Lean 4 problem, where it generates Lean 3 syntax (e.g. begin).

Prompt: Write and prove the following in Lean 4: the sum of the first n positive integers is $n \mapsto \frac{n(n+1)}{2}$. Begin with "import Mathlib". Surround your Lean 4 code with “lean4 and “

```

/-- The sum of the first n positive integers (i.e.  $0 + 1 + \dots + n$ ) is  $n * (n + 1) / 2$ . -/
theorem sum_first_n (n : ℕ) : (range (n + 1)).sum id = n * (n + 1) / 2 :=
begin
  induction n with n ih,
  { -- Base case: when  $n = 0$ , we have  $\text{range } (0 + 1) = \{0\}$  and so the sum is 0, which matches  $0 * \mapsto (0 + 1) / 2$ .
    simp [range],
  },
  { -- Inductive step: assume the formula holds for  $n$ , and prove it for  $n + 1$ .
    -- First, note that  $\text{range } (n + 2) = \text{insert } (n + 1) (\text{range } (n + 1))$ .
    rw [range_succ],
    -- Since  $(n + 1) \notin \text{range } (n + 1)$ , we can express the sum as:
    rw [sum_insert (not_mem_range_self (n + 1))],
    -- Substitute the induction hypothesis:
    rw [ih],
    -- We now have to verify:
    --  $n * (n + 1) / 2 + (n + 1) = (n + 1) * (n + 2) / 2$ .
    -- This equality follows by algebra (ring calculation).
    ring,
  }
end

```

3.9. High Logical Complexity and OOD Domains

Tasks such as writing highly concurrent code or discovering performance optimizations have a high logical complexity, often proving difficult for even the best human coders. Similar to solving research-level math problems, these out-of-distribution domains are very hard for LLMs.

Potential solutions: [4.2.1](#)

Some programming tasks are challenging for even the best human programmers, requiring approaches with a very high logical complexity. Examples of tasks that fall into this category include superoptimizing programs, discovering attacks for purportedly secure code, writing performant compilers, optimizing GPU kernels^[37], and writing very error-prone and very technical code.

Example. Synthesis of Sorting Kernels: An example of an out-of-distribution domain is synthesizing fast assembly code for sorting kernels. In 2023, AlphaDev^[112] used reinforcement learning to find a SoTA kernel for sorting length 3-5 arrays. While this appeared to be a superhuman performance, shortly after,^[113] hand-wrote a kernel shorter and faster than the one found by AlphaDev. Later,^[114] developed an algorithm based on enumeration and intelligent heuristic-based sampling that beat both of these. In addition, the algorithm ran faster than AlphaDev by two orders of magnitude. In this case, while AI was able to achieve an impressive performance, humans were able to discover better algorithms.

Example: Verifying File System Properties: In formal verification, when working with new domains, it is necessary to devise new theories to faithfully represent desired properties. For example, FSCQ is a formally certified crash-proof file system with the provable guarantee that under any sequence of crashes followed by reboots, FSCQ will recover the file system correctly without losing data^[115]. In this domain, one challenge is that proving safety cannot be done at the source code level—because instructions are not atomic, data may be lost if the crash occurs within a non-atomic instruction. Instead, a new logic known as the Crash Hoare logic (CHL) needed to be developed, and constructs representing a crash condition and recovery procedure needed to be described. Constructing a logic like this would be very difficult for AI systems.

Limits of Symbolic Techniques: When it comes to applying symbolic techniques to these tasks, there are a few limiting factors that make them difficult to tackle. First, for synthesis-style tasks, the search space can be very large. Deductive and rewrite-based synthesis techniques are unable to explore a majority of the search space. Second, verifiers can be limited in power, such as when dealing with properties in concurrency or weak memory models. Third, many domains lack clean models to reason about properties, such as dealing with memory bandwidth in GPU kernels.

Because they are hard for humans, these tasks are very rarely in the training data of today’s language models. They have unique, domain-specific, challenges that making generalizing from existing data difficult. For these problems, language models rely heavily on feedback-driven search algorithms^[112], and it can be difficult to navigate the search space effectively. In addition, many of these tasks lack feedback mechanisms, which is crucial for AI to pick up learning signals. When designing a complex algorithm or data structure, it is often hard to know if you are on the right track until you get to the correct result.

When writing code for a large multithreaded operation, it may be hard to know if the algorithm has concurrency issues until all the parts are fully fleshed out. Without feedback, incremental improvement is nearly impossible.

4. Paths Forward

4.1. Data Collection

One bottleneck in the development of AI for SWE in the open-source community is the lack of access to fine-grained and high-quality code data. In Sec. 4.1.1, we discuss how automated techniques can mitigate this by augmenting existing programs with symbolic information and generating synthetic data with symbolic verifiers. However, there are other crucial signals in programming that can be hard to automate. We envision that a large community-based coding data curation effort will be very impactful. In Sec. 4.1.2, we discuss examples of datasets that such a community could create that would unlock new capabilities in code LLMs.

Challenges addressed: 3.1, 3.3, 3.6

4.1.1. Automatic Data Curation

Augmenting Data with Program Information: One challenge in enabling LLMs to develop a world model of code is that programs are often treated like text: as tokens with no semantic information. However, modern programming tools allow us to extract rich semantic and structural information about code. By leveraging these tools, we can augment training datasets with detailed annotations describing various properties of programs. We hypothesize that this augmentation will significantly improve a model's understanding of code, leading to better generalization and stronger coding capabilities. Information can include:

- Static analysis: the syntactic structure of a program (abstract syntax trees, control flow graphs), information about the type of each variable, data flow analysis (reachability, liveness analysis)
- Program instrumentation: memory consumption, runtime analysis, aliasing, and code coverage (like statement or branch coverage)
- Dynamic analysis: program states at various points in the program, call stacks, dynamically curated properties (often relies on instrumentation)
- Formal verification: concurrency analysis, program invariants, loop invariants, memory safety

There have been a few examples of this in the literature: [37] leverage profiler feedback to improve GPU kernel generation, [100][116][99] incorporate execution trace information, [117] train with program invariants, GraphCodeBERT [118] incorporate data flow information, and [119] train on a dataset of performance-improving edits.

High-quality, Verifiable Synthetic Data: The advantage of code is it is possible to achieve strong, verifiable feedback with test cases, program execution engines, and other symbolic tools. This makes high-quality synthetic data generation viable, as it is possible to generate a large batch of data and filter out low-quality samples. For example, to generate code with interesting program invariants, we can sample a large batch of programs, run an invariant detection engine, and retain only programs with interesting invariants. While synthetic data in code has mostly been at the function-level scope, there are no fundamental bottlenecks to expanding to larger scopes. As code is quite compositional, individual building blocks can be combined to generate complex synthetic data at the repository-level scope, which can be very helpful in both training and evaluation.

While the importance of having high-quality data vs. high quantities of data is debated, using verified data has proven to be useful. For example, [120] shows that simply removing bugs in existing datasets such as TACO [121] can lead to significant boosts. KodCode [57] also showed that fine-tuning on verified synthetic data also leads to significant improvements. However, these works work with programs at the function-level scope with low to medium logical complexity, and we imagine that general SWE abilities can improve with synthetic data across scopes and logical complexities.

In DSLs, where programs can be cleanly described with semantics and rewrite rules, one can symbolically generate programs with desired properties via sampling, drawing on enumeration techniques from program synthesis[122]. This technique has been successfully applied to make considerable progress in difficult reasoning tasks such as ARC-AGI[123] and math olympiad problems[124][125][126].

4.1.2. Human-Centric Data Curation

Below, we list three classes of human-annotated data that would be invaluable for the next generation of coding LLMs.

Fine-Grained Data of the Developmental Process: Many code LMs are trained on datasets such as the Stack[127][128], consisting of trillions of tokens sourced from GitHub. However, training on raw GitHub tokens omits many crucial human signals in the process of software development. For example,

companies such as Google rely on internally captured logs of high-quality SWE data. This includes "fine-grained code edits, build outcomes, edits to resolve build issues, code copy-paste actions, fixes of pasted code, code reviews, edits to fix reviewer issues, and change submissions to a repository"^[129]. Similarly, Meta and GitHub Copilot use telemetry with their AI coding assistants to track and leverage signals from AI-generated code^{[130][131]}. These tools, along with coding IDEs like Cursor, could provide a treasure trove of reward data for RL-based methods. With direct access to the full history and evolution of a codebase, they can track which suggestions are adopted over time. However, collecting data from human usage also raises critical privacy and intellectual property concerns.

Data for Diverse SWE Tasks: Most of today's code LLM training recipes still focus primarily on code generation because large-scale datasets are mostly in a continuous, tokenized format. However, as described in Sec. 2), there are many tasks involved in software engineering which models lack exposure to. Training on a broader set of tasks would also incentivize models to learn *general* capabilities of programs beyond just generation (e.g. a better understanding of program semantics). As initial evidence, ^[132] find that training models on input-output prediction data leads to consistent improvements on reasoning tasks.

The lack of high-quality data on these tasks makes it hard to train on them. It can also be hard to automatically curate them on GitHub. For example, for code refactoring (Sec. 2.2.1), we need paired repositories before and after refactoring, ideally with the refactoring changes described. While some signal such as commit messages and version releases can be used, many repositories lack clean commit histories and releases conflate many features at once. Therefore, to mitigate this, we envision large community-based efforts curating task-specific data on these diverse challenges.

Human-Centric Data: Code LLMs are typically trained and evaluated on carefully curated datasets with clear instructions and verifiable test cases. However, as discussed in Sec. 3.3, these models are often deployed in real-world scenarios where users provide vague specifications or incomplete requirements in their queries. Collecting human-centric data that reflects real-world model usage is a promising approach to bridging the gap between model development and deployment. Recent efforts, such as Copilot Arena^[50] and WebDev Arena, have explored gamified arenas to gather data on human preferences, offering an alternative to purposefully curated datasets. However, such data collection methods may introduce noise, and arena-style approaches are not well-suited for long-horizon, interactive tasks. One potential approach is to leverage existing coding tools and environments, such as developing plugins for GitHub Copilot^[133] or open-source IDEs, to capture real-world interactions. Unlike

static datasets, human-centric data can also be collected encompassing diverse interaction modalities, such as users providing sketches to AI coding systems for web development^[134]. As AI coding systems continue to emerge and evolve, launching data initiatives focused on human-centric SWE data is also a crucial direction for advancing human-AI collaboration in software development.

4.2. Training

4.2.1. Environment Design for Code RL

Reinforcement Learning from Verifiable Rewards (RLVR) (Lambert et al., 2025) has emerged as a powerful paradigm in math and coding domains where model outputs can be evaluated against a ground truth outcome such as exact match and passing a set of unit-tests. Towards this direction, promising avenues include collecting executable codebase environments, sourcing task prompts/rewards from GitHub, and designing non-execution based rewards based on program syntax and semantics.

Challenges addressed: 3.4, 3.9

Collecting executable codebases: In recent months, RLVR has seen success in solving algorithmic programming problems through DeepSeek-R1^[135] and OpenAI o1. Recently, on SWE-Bench, SWE-RL^[136] use RL on a rule-based reward to improve performance on SWE-Bench. We find it promising to continue scaling the RL approach to problems collected from real-world software engineering repositories. Towards this, we believe that collecting execution-assisted gym-like reinforcement-learning environments will lead to further performance improvements. These environments can be used further to improve reasoning skills, environment-interaction capabilities and tool usage.

Several prior works^{[86][137][138][139]} curate executable environments for programming agents by supporting CI/heuristic-based repository installations. However, these works are at a relatively small scale and limited in scope, offering only a few thousand tasks from a maximum of a thousand repositories and more importantly, limited to the Python language. Scaling this up significantly requires solving several research and engineering problems. First, installing arbitrary repositories from Github, even using CI is challenging and we require smarter solutions potentially involving LLM-based installation agents. Next, setting up execution infrastructure would require storing installed repository images in something akin to docker for efficient storage and fast container startup times^[140]. Notably, combined docker images can grow massively large and often grow at hundreds of gigabytes even at a

modest scale of a few hundred repositories. They require engineering support for efficient storage and serving of such images.

Sourcing task prompts and rewards: Beyond environments, performing large-scale reinforcement learning would require collecting diverse challenging problems with an appropriate way to compute the rewards. These task prompts can be collected from Github^[137] or generated synthetically from problems on Github. Moreover, assuming access to many executable repositories, we can source various end-to-end problems for tasks beyond bug-fixing such as optimization, fuzzing, etc. Access to pre-existing or generated test cases allows for measuring correctness and providing rewards.

However, we envision many practical challenges to remain. For example, longer-horizon tasks are usually more ambiguous and approaches may require multi-turn interactions beyond autonomous coding agents. This would pose a considerable challenge during reinforcement learning where ambiguity resolution might need to be modeled in the reinforcement learning process itself. We elaborate on human collaboration further in Section 4.2.3. Reward hacking^[141] poses another challenge as we build more real-world coding challenges. Test cases often suffer from coverage issues and can grade correct solutions as incorrect. For example, ^{[142][143]} identified that models attempt to bypass or cheat against the testing harness when optimized using reinforcement learning.

Rewards without execution: As setting up execution environments can lead to considerable overhead, another potential strategy is to use proxy metrics and trained language models to judge correctness. This was common in the pre-LLM era, researchers often used BLEU/CodeBLEU^{[144][145]} and BERTScore/CodeBERTScore^{[146][147]} to assess correctness of text and code. In code, semantic and structural properties can be used to improve similarity metrics. Two examples of this are Dolos^[148], an AST-aware plagiarism detector, and difflib.SequenceMatcher, which can be used to compute the similarity between two patches^{[136][149]}. Beyond rule-based rewards, LLMs-as-a-judge approaches can also be used as reward functions, possibly in conjunction with other execution-based or execution-free approaches.

4.2.2. Adapting to Specialized and Quickly Changing Codebases

Low-resource languages (Sec. 3.7), custom APIs, library version updates (Sec. 3.8), large codebases (Sec. 3.5), and custom coding styles all surface the fact that code LMs struggle to adapt to unseen specialized contexts. Customization can be achieved through test-time training, keeping specialized information in an information bank. A cheaper and alternative approach to test-time training is to apply prompt and prefix tuning, where codebase-specific embeddings are learned and applied depending on the context.

Challenges addressed: 3.7, 3.8

Test-time training (TTT) to custom codebases: TTT is the recent paradigm of adapting to a specific problem instance by training on a narrow set of in-distribution examples^{[150][151]}. This can be used when working in a low-resource context, for example training on a specific codebase, new domain, or unseen API. One challenge in this setting is customizing the model to the particular codebase while retaining general coding knowledge, potentially by using algorithms that can induce controllable forgetting^[104]. To get data in specialized contexts, we envision two mitigation strategies: generating synthetic data and collecting trajectories. In-distribution synthetic data can be generated in large quantities and then filtered and annotated with symbolic (e.g. compiler) information to gain a more global understanding of the current environment and setting. To gather agentic trajectories, we can keep track of previous model attempts and failures to learn from past successes and avoid making repeated mistakes. This will steer the model closer to the desired distribution—for example, to generate code in the specified version of libraries being used in the current context.

Keeping an information bank of code information: For library and versioning issues, retrieval (Sec. 4.3.1) can be very effective for preventing hallucinations of wrong versions of libraries, which can inherently lead to better synthetic data and agentic trajectories. During the TTT process, we can also keep a large growing memory bank of code, documentation, synthetic code, and agentic trajectories in the specialized context. Retrieving from the memory bank would improve the success of generating code, which can then be augmented to the memory bank, and so on, continuously increasing the amount of data and knowledge available.

Prompt and prefix tuning for specialized code contexts: One issue that makes it difficult to continuously keep up with library updates is that doing full finetuning every time something changes is very expensive. Because only a small amount of knowledge needs to be learned compared to that of the pre-

trained model, we believe less expensive approaches such as prompt tuning^[152] and prefix-tuning^[153] could suffice. Both these methods append a set of learned task-specific vectors to the input sequence in order to model a specified context, though prompt tuning only modifies the input and prefix-tuning modifies the input at each layer. These methods have also been shown to have good OOD performance, and we believe they present a promising approach to dealing with multiple library versions. A separate prompt/prefix can be trained for each version and then applied according to the context. When an API has new updates, the prompt/prefix can then be cheaply re-tuned to reflect the new updates without undergoing full fine-tuning. This approach also applies to adhering to specific coding styles, where codebase-specific prompts/prefixes can also be learned.

Learning on the fly: When humans are faced with a task they have never seen before, they are often able to draw from past experiences and quickly adapt and generalize to the new domain. This is one of the big unsolved challenges of today's LLMs: given an OOD coding task, how can models get up to speed and productively work on the task with few samples? On toy domains, an example of this is DreamCoder^[81], a system that learns to solve problems by writing programs and automatically discovering domain concepts and composing concepts together. Designing such approaches for more practical applications is an exciting research direction that will have drastic implications for coding and reasoning.

4.2.3. Training Code LLMs to Collaborate with Humans

Training the next generation of code LLMs needs to account for human-AI collaboration, as these models will likely be deployed in ambiguous and interactive scenarios. We highlight two key directions for improving collaboration: First, learning to leverage specifications beyond natural language through formal methods and user-specified tests can mitigate vague specifications. Second, improving uncertainty quantification and proactive communication through post-training has the potential to prevent hallucination and misalignment.

Challenges addressed: 3.3

Learning to Leverage Specifications Beyond Natural Language: As discussed in Section 3.3, while natural language prompts offer intuitive and flexible ways to express requirements, they often suffer from ambiguity and incompleteness. One direction to address this limitation is to train models to leverage enhanced specifications with more precise and verifiable representations, such as formal specifications and test-based specifications.

Formal specifications: To mitigate underspecification issues, one solution is to develop systems that can translate user intent into formal specifications^{[154][155]}. While current autoformalization approaches face challenges in accurately capturing user intent (see example below), we envision next-generation systems that will iteratively refine formal specifications through interactive verification with human feedback. These systems would present intermediate formalizations in accessible notation, enabling non-expert users to verify correctness before code generation.

Example: Incomplete specification in Verus: Here, we show a failure mode of LLMs when writing specifications and proofs in Verus. The LLM is asked to write the `ensures` postcondition clause for a ring buffer enqueue function^a. Here, the postcondition is incomplete: it does not check, for example, that the original elements were maintained in the ring buffer.

```
fn enqueue(&mut self, val: T) -> (ret: bool)

  ensures
    ret == !old(self).is_full(),
    self.inv(),
    if ret { self.view() == old(self).view().push(val) } else { self.view() == old(self).
    ↪ view() }

{
  if self.is_full() {
    false
  } else {
    self.ring.set(self.tail, val);
    self.tail = (self.tail + 1) % self.ring.len();
    true
  }
}
```

^aFull example [here](#)

Tests as specifications: Another approach to specify software behavior is through tests. These range from input-output examples and assertions to property-based tests. However, in practice, hand-crafted test suites are often incomplete, failing to capture the full intended behavior, particularly edge cases. This can lead to misalignment, where AI-generated code passes tests but does not genuinely meet functional requirements, potentially misleading users. Moving forward, a direction is training models to generate high-quality test cases based on the user's initial query, ensuring more comprehensive specification coverage.

Example: For instance, in a release of [AI CUDA Engineer](#) by Sakana AI, an AI-generated CUDA kernel for lower triangular matrix multiplication—purportedly achieving significant speedups—was later found to exploit out-of-bounds memory access to bypass correctness checks^a. Advancing research on frameworks that facilitate test generation and automated adversarial testing represents an important direction.

^a The full LLM-generated kernel code can be found in Listing 3, pg. 46–47 of [\[156\]](#).

Learning to Quantify Uncertainty and Communicate Proactively: As AI coding systems are increasingly deployed to complex software engineering tasks, they encounter more ambiguous and uncertain scenarios compared to traditional benchmarks for coding models. Ideally, in such situations, these systems should proactively communicate with users to clarify tasks and acknowledge its own limitations rather than becoming stuck in endless failure loops or generating buggy code. A key challenge is enabling models to distinguish between well-specified and ambiguous instructions while quantifying uncertainty in a robust manner. While early studies, such as [\[157\]](#) and the example below, demonstrate that interactive LLMs can improve performance through clarification-seeking behavior, current models still struggle with uncertainty estimation. Equipping models with the ability to quantify uncertainty will likely require incorporating corresponding reasoning data into the post-training stage.

Besides uncertainty quantification, [\[158\]](#) identify communication as a primary challenge in human-agent collaboration, highlighting the need for improving models’ proactive communication capability. Current models often fail to ask meaningful questions when user input is ambiguous or insufficient, and they struggle to provide progress updates or verify plans in interactive settings. Enhancing models’ proactive communication abilities requires innovative approaches to reward behaviors that yield benefits over multiple steps. Since communication with users does not immediately resolve the task at hand but may improve long-term outcomes, effective strategies must account for delayed rewards in training.

Example: Discussion Helps Coding Agents Resolve Github Issues: In SWE-bench^[20] pydata_xarray-4750, the original issue description requests limiting the number of data rows displayed in repr. While it suggests a maximum of 25 rows, it does not specify whether this number should be configurable—a key requirement that emerged during the issue discussion. When SWE-Agent^[66], powered by GPT-4o, uses only the issue description as the problem statement, it generates a function that hardcodes the maximum at 25, causing the solution to fail the test. However, incorporating the issue discussion allows the agent to produce a correct, test-passing implementation (see Listing 2). This suggests that enabling coding agents to engage in discussions with users could potentially improve the issue solving rate.

```
# Use only the issue description
def _mapping_repr(mapping, title, summarizer, col_width=None):
    if col_width is None:
        col_width = _calculate_col_width(mapping)
        summary = [f"{title}:"]
    if not mapping:
        summary += [EMPTY_REPR]
    else:
        items = list(mapping.items())
        if len(items) > 25: # Show first 12 and last 12 items
            first = items[:12]
            last = items[-12:]
            summary += [summarizer(k, v, col_width) for k, v in first]
            summary += ["    ..."]
            summary += [summarizer(k, v, col_width) for k, v in last]
        else:
            summary += [summarizer(k, v, col_width) for k, v in items]
    return "\n".join(summary)

# Use the issue description and discussion
def _mapping_repr(mapping, title, summarizer, col_width=None):
    col_width = _calculate_col_width(mapping)
    summary = [f"{title}:"]
    if mapping:
        items = list(mapping.items())
        if title == "Data variables" and len(items) > OPTIONS["display_max_rows"]:
            # Show first and last variables if there are too many
            first_n = OPTIONS["display_max_rows"] // 2
            last_n = OPTIONS["display_max_rows"] - first_n
            selected_items = items[:first_n] + [(".", "...")] + items[-last_n:]
        else:
            selected_items = items
        summary += [summarizer(k, v, col_width) if k != "..." else "    ..."
                    for k, v in selected_items]
    else:
        summary += [EMPTY_REPR]
    return "\n".join(summary)
```

Listing 2. SWE-Agent improves when incorporating issue discussions

4.3. Inference Time Approaches

4.3.1. Semantic-Aware Embeddings and Retrieval

In contrast to text, embeddings for code should incorporate execution and semantic information, improving retrieval. RAG benefits from both context-aware retrievals and explicit training on how to use them, enhancing code reuse across languages and APIs. Beyond static retrieval, AI agents could also dynamically navigate codebases using command-line tools and IDE functions.

Challenges addressed: 3.5

Semantic and execution aware code embeddings: When training LLMs, code is often treated as pure tokens (just like text) rather than explicitly incorporating code-specific information such as program execution and semantics. As a result, code that is close in embedding space is more often syntactically similar than semantically similar^{[92][159]}, and there are few reliable methods today to retrieve semantically similar code. However, before the LLM era, there were a variety of efforts to incorporate code properties when training embeddings. For example,^[160] train neural modules to represent program operations, leading to compositional program representations that encode the semantics of the underlying programming language. Many other works^{[161][162][163]} attempt to learn execution-aware latent representations for partial and full programs, taking semantics into account.

We speculate that incorporating these techniques to train models to have better and more semantically aware representations may lead to models with a more general understanding of code (Sec. 3.6). For example, if correct and buggy programs could hypothetically be separated in embedding space, then models could be steered away from the incorrect program space. While such a clean separation might not be possible, we believe that training embeddings to have interesting semantic properties is worth exploring.

Better retrieval-augmented code generation: When retrieval-augmented language models were first introduced, they often relied on training the retriever and language model jointly, as in FiD^[164], RETRO^[165], and Atlas^[166]. As language models increased in size, the field shifted to a black-box setting^[167], where the retrieval module is tuned independently to adapt to the pretrained black-box LLM. This setting is much more cost-effective, but the language model is not explicitly trained on how to use its retrievals.

The black-box setting is ideal for challenges such as low-resource languages or specialized contexts. In these situations, the model has not seen enough training data to fully grasp the context, and the challenge is often syntactic rather than algorithmic. For example, when adapting to a domain or a codebase where the relevant API functionality or code style, retrievals can be very instructive. When using APIs with multiple versions, providing retrievals in the correct version can inform the model of how to use the API. When writing code in a completely new language, showing examples of for loops and while statements will teach the model the syntax of these constructs. Retrievals should be diverse and given in multiple forms, including documentation, function definitions of APIs that are used, and example use cases of target functions.

In many other cases, however, we believe that a black-box setting is insufficient. As described in Sec. 3.5, there are two challenges: 1) knowing what to retrieve and 2) using the retrieval. The first challenge relies on retrieving relevant examples, both syntactically and semantically. We believe that having more semantically aware embeddings, as mentioned above, will drastically improve this. For example, embeddings can be trained contrastively to minimize the distance between semantically similar programs. Another potential direction is to consider a diverse set of potential retrievals and then train the retriever to prefer samples that help during generation, as in Atlas^[166].

The second challenge, using the retrieval, is a code reuse task, which requires complex reasoning and code understanding. Algorithms provided in retrievals may often need to be modified and adapted significantly to adapt to the current setting. An example of this might be writing a C++ version of a shortest path algorithm when the retrieval is a Java version, a translation task that models may not have been trained for explicitly. Long chunks of retrieved documentation may need to be understood precisely so that correct hyperparameters and flags can be used. Yet, in a black-box setting, models have not been explicitly trained to leverage this information. Therefore, just as training on incorrect-correct code pairs can improve program repair, we believe that direct training can be very beneficial for code reuse and retrieval-augmented generation. Execution information could also be useful, as code reuse often requires understanding the situation well enough to identify subtle differences between the context of the retrieved code and the current context.

Retrieving via code navigation on the fly: Standard retrieval-augmented methods keep a large retrieval index containing millions of embeddings, which can require a high one-time cost to create. As the codebase evolves, these embeddings may also need to be continuously updated. Instead of keeping track of embeddings, another approach is to find retrievals on the fly by navigating the codebase. We can

imagine an agent that learns to use command line functions such as `cd`, `ls`, and `grep`, as well as IDE functions such as jumping to function definitions or finding all references of a function. Static analysis tools can also be paired with the agent to improve code navigation, such as providing the abstract syntax tree (AST) or file structure of a codebase.

4.3.2. Integration with SWE Development Frameworks

Integrating AI with SWE development frameworks is critical for practical applications and impact on developer workflows. While software development is inherently integrated with tools, workflows, scaffolding, and meta-code, these are often absent from source code and scarce in AI training data. Ensuring that AI deeply understands software deployment beyond code editing is crucial, as writing code is only a small part of the development cycle. These can include automated reviews, deployment risk assessments, and documentation generation. We can also fine-tune LLMs to recognize and avoid known software anti-patterns such as CWEs.

Challenges addressed: 3.4, 3.5

Incorporating AI into the CI/CD process: In continuous integration and continuous deployment (CI/CD), automated pipelines are the backbone for building, testing, and deploying code changes. CI/CD accelerates feedback cycles and minimizes integration issues. AI offers several integration points within CI/CD. AI-powered code review tools can be incorporated into CI pipelines to automatically identify and flag style violations, potential security vulnerabilities, and code smells before human reviewers are involved. Furthermore, AI can provide intelligent deployment risk assessments. By analyzing code changes, test outcomes, and historical deployment data, AI can predict the likelihood of deployment issues, informing decisions about whether to proceed with automated deployment or mandate manual verification steps. Finally, AI can automate the generation of release notes by summarizing commit messages, issue tracker data, and relevant code modifications within the CI/CD process.

Steering away from software anti-patterns: In software engineering, certain anti-patterns frequently lead to bugs. For example, common weakness enumeration (CWE) is a categorization of software and hardware weaknesses often leading to vulnerabilities. Because publicly available GitHub code often contains code with anti-patterns, bugs, and CWE vulnerabilities, LLMs often write code susceptible to these issues^{[168][169]}. We hypothesize that explicitly steering models against these vulnerabilities will lead to more secure and correct code. One way to do this is to collect a large number of program samples

violating each CWE (either synthetically or on GitHub) and then use these samples as negative signal during further supervised fine-tuning or RL stages.

4.3.3. Incorporating SWE Tools

Software engineers integrate a variety of domain-specific tools when writing code. By repeatedly interacting with tools in an RL-style manner, AI can develop the ability to do the same. Beyond tool use, using neurosymbolic approaches such as incorporating program analysis and type-checking can also help enhance LLM capabilities.

Challenges addressed: 3.2

Learning to use SWE Tools: As mentioned in Sec. 3.2, we believe SWE agents should understand the intricacies of programming tools and be able to autonomously invoke them as needed. There are three skills to learn: which tool to use, how to use the tool, and how to incorporate the results of the tool. Similar to how models learn to play complicated games, we believe that intelligent tool integration can be learned through repeated interactions with the tool in a RL-style manner. One way we envision this is as follows: first, the interface of the tool must be precisely specified. Next, data containing repeated interactions from the tool (with varying degrees of success) should be collected. Finally, multiple rounds of RL and expert iteration can be done to improve understanding of the tool and learn from misuses.

Evidence that learning higher-level strategies might be possible is that through test-time techniques, OpenAI’s o3 model learned to write brute-force solutions to verify the correctness of more complicated solutions^[170]. We envision that after learning to use tools, AI coding agents can autonomously invoke tools as needed to improve its overall world model of the code and hence its software engineering capabilities.

Neurosymbolic Approaches: Code is a unique domain because there is a vast body of techniques from programming languages (PL) research to build off of, but the majority of AI for code research today does not leverage the symbolic properties of code. Some of these PL techniques are as follows: abstract interpretation^[171] is a technique to compute over-approximations of program state in order to prove the soundness of program properties at points in the code. Concolic testing^{[172][173]} finds bugs in software by combining concrete and symbolic execution. Model checking^[174] is a way to prove properties of execution traces via temporal logic. Linting and type-checking^[175] provide a static check to ensure that variables, expressions, and functions adhere to a programming language’s rules. Finally, many other

program analysis algorithms leveraging these tools have been designed to prevent bugs and ensure code correctness properties.

Traditional PL approaches have a few common shortcomings, which overlap with some of the issues mentioned in Sec. 2.6. First, they often require very complete and precise specifications. Many tools need to have specifications for all library functions, need to specialize to a precise version of the language, and need to specialize to the build system. Second, there is often a high computational cost due to the large search space. Third, there can be many false positives due to the limitations of the tool. We believe that deeply integrating these symbolic tools with LLMs can partially mitigate these challenges.

We provide a few examples of this potential integration. When generating code, program analysis techniques could be applied on shorter snippets of AI-generated code to surface potential bugs or prove properties of the generated code. To improve general code understanding, LLMs can be trained with information about program structure such as abstract syntax trees^[176]. When debugging a large codebase, when the scale is too large to directly apply PL techniques, AI could be first used to narrow down potentially problematic sections of the code which are then handed off to PL tools for debugging. During code generation in DSLs, LLMs can leverage the grammar of the programming language to do constrained decoding^{[177][178][179]} to mitigate syntactic errors. During code refactoring, abstract interpretation and static analysis can be used to identify whether new errors have been introduced and preemptively cut off unpromising search paths.

Deductive Synthesis and Intermediate Languages: Early program synthesis relied on *deductive synthesis* approaches^[180], where programmers would write a clean simple implementation and then apply transformation rules to convert it into a more efficient one. The appeal of deductive approaches is that because these rewrite rules are semantics preserving, there is a correct-by-construction guarantee. One success story of deductive synthesis is Spiral^[181], a DSL for signal processing kernels that takes advantage of domain-specific transformation rules to produce implementations beating expert hand-crafted code. Another example is Halide^[182], a DSL for high-performance image and array processing code. Due to the difficulty of writing optimized code, humans generally opt for writing code in these intermediate DSLs, and we find it promising for LLMs to do the same.

Example. LLM-aided Compilation for Tensor Accelerators: As an example,^[183] consider the task of generating highly optimized, hardware-efficient code for a tensor accelerator from a high-level specification of a computation (e.g. C++ code). Their pipeline works in two steps: first, the high-level specification is translated to a DSL. Then, the DSL code is symbolically compiled to hardware-specific instructions. The LLM is also used to optimize the DSL code via a cost model driven search, where it suggests rewrites and scheduling operations (e.g. loop reordering) that guarantee semantic equivalence.

4.3.4. Scaffolding Human Supervision

At inference time, most machine-generated code will be presented to humans in a format shaped by the human-AI interface design. Since AI may be responsible for generating the majority of the code within a human-AI team, it is important to ensure human control and oversight. By scaffolding human supervision with techniques like summarization and interactive verification, we could potentially improve trust in AI-generated code.

Challenges addressed: 3.3

At inference time, most machine-generated code will be presented to humans in a format shaped by the human-AI interface design. Since AI may be responsible for generating the majority of the code within a human-AI team, it is important to ensure human control and oversight. By scaffolding human supervision with techniques like summarization and interactive verification, we could potentially improve trust in AI-generated code.

Challenges addressed: 3.3 Once code LLMs are deployed for inference, it is crucial to scaffold human supervision of AI-generated code. This goes beyond merely enhancing the accuracy of AI-generated code, as humans often still need to make the final decision on whether to accept the code or understand it for future integration and maintenance. A study on Github Copilot usage^[184] revealed that programmers tend to allocate less visual attention to AI-generated code. While one solution is to train humans to better identify issues in AI-generated code^[185], a more desirable approach is to design AI systems that scaffold human supervision, reducing their cognitive load when reviewing generated code.

One way to achieve this is by enriching AI-generated content with additional contextual information. Modern LLM chatbots now routinely generate text with citations for knowledge-intensive queries. In Collaborative STORM^[186], researchers demonstrated that dynamically presenting hierarchical “mind

maps” alongside the actual collected information significantly enhanced human-AI collaboration, particularly in long sessions. In software engineering specifically,^[187] highlighted the benefits of high-quality source code summarization in aiding software developers in understanding and maintaining machine-generated code. Second, interactive approaches can also enhance supervision. One example is *Live Programming*^[188], a continuous display of the runtime values of a program, as a means of lowering the cost of validating AI-generated code. However, these existing studies are largely limited to specific programming languages and small codebases. Finally, improving the readability and interpretability of AI-generated code itself presents a promising direction. For example,^[189] showed that modeling program synthesis as rational communication improved end-user interpretation and subsequent communication of code. Expanding on these ideas, future research should prioritize human interpretability in the design and optimization of AI coding systems, fostering greater trust and control in AI-assisted software development.

5. Limitations

We identify a few limitations below:

Speculative nature of future work: The ideas we list in the future work section are opinionated directions we believe have a high chance of success. Many draw upon insights from related work in the literature, but many lack strong and concrete evidence. We encourage further research validating or disproving the effectiveness of these ideas.

Limited scope of future work: We also do not include any novel moonshot ideas, and many of the directions we propose have their roots in existing code LLM literature. Our future work section is also relatively general and applies holistically to AI for code. However, the field has many tasks and challenges that can benefit from using domain-specific knowledge and insights, and we do not touch on these. Finally, this paper is written by people primarily in the academic community, who may not know the details of cutting-edge methods employed in frontier industry labs. We cater this paper towards areas we have more expertise in, and thus leave out many promising directions such as novel architectures.

Focus towards code-specific challenges: In this paper, we mostly focus on code-specific challenges and techniques. However, there are many techniques that apply to general LLM reasoning and development that could be directly applied to code. We believe many of these methods can be used in synergy with code-specific techniques.

Quickly changing nature of the field: The field of LLM for software engineering is progressing very rapidly, with new innovations released weekly. It is possible that a reader reading this paper a few months down the line will find that several of the mentioned challenges will have been partially or entirely resolved.

6. Conclusion

In this position paper, we have identified key tasks at the heart of AI for software engineering as well as a set of three measures to classify different realizations of these tasks. We have also highlighted critical cross-cutting challenges that permeate throughout many tasks. Finally, to drive progress in AI for code, we've pinpointed a set of exciting and promising research directions for alleviating these challenges and advancing AI towards being a more capable software engineer. We hope this work provides valuable insights about the current landscape of AI for software engineering and encourages future research in these directions. By building on these insights, we are optimistic that the community can work toward developing AI-driven solutions that better support software engineers in real-world settings.

Appendix A. Survey of Related Work: Tasks in AI Software Engineering

In this section, we briefly survey some of the relevant works for each of the tasks we mention in Sec. 2. These works are by no means complete, and we encourage the reader to check out the survey works mentioned in the introduction and in this section for further references.

A.1. Code Generation

Code Completion: Completion typically happens in conjunction with live programming or within an IDE, helping developers write code faster by suggesting relevant continuations. Traditional code completion systems rely heavily on syntactic and type-aware models (e.g., AST-based models), but recent advances leverage LLMs trained on code corpora to offer semantically rich and context-aware suggestions, naturally following the next-token prediction task in language modeling^[190]. Tools like GitHub Copilot and Codex exemplify this trend^[15], and are followed by commercial tools such as Cursor⁴ and Tabnine⁵. Recent advances in context-aware^[191], grammar-aligned^[192], and constraint-based decoding^[193] have improved the quality of local completions, particularly for shorter code snippets. For longer code

snippets, the typical task formulation is method implementation synthesis given a function signature. This setup is commonly evaluated using benchmarks such as MBPP^[16] and HumanEval^[15].

Natural Language to Code Generation: Translating natural language into code has long been a central challenge in AI for programming. Early attempts at code generation involved semantic parsing^{[194][195]}, where natural language is translated into logical forms or domain-specific languages. A prominent example is SQL query synthesis from natural language questions, as seen in systems like Seq2SQL^[196] and Spider^[197], where the target language is constrained, small, and domain-specific. Recent work demonstrates that large language models (LLMs) can generalize to general-purpose programming languages, enabling the generation of larger and more complex code snippets^[198]. When applied to code completion, users often begin with natural language instructions in the form of comments, which LLMs use as context for code synthesis. Beyond function-level code generation^{[16][15]}, recent work has extended to class-level generation^[199], which targets classes in object-oriented programming, and even project-level code generation^{[200][201]}, which involves generating or completing entire multi-file codebases.

Multimodal Code Generation: While text can describe most cases of code generation, certain instructions are better defined visually. For example, in graphics applications, visual context such as a trajectory or a 3D model is essential to synthesize the correct code. Demonstrations of GPT-4's multimodal capabilities have shown that models can generate functional webpage code directly from paper sketches, translating visual layouts into HTML and CSS^[202]. LogoMotion^[203] explores visually grounded code synthesis for animations and motion graphics in JavaScript. The system leverages vision-language models (VLMs) to incorporate both visual inputs and user instructions, enabling code generation that aligns with spatial and temporal visual cues. Other works, such as SynthesizeCAD^[204] and SGP-Bench^[205], explore how LLMs can interface with visual and 3D modalities by generating code in languages like SVG and CAD.

Code Generation in Low-Resource Languages: As discussed in Sec. 3.7, one major challenge is writing code in low-adoption general purposed language and domain specific languages (DSLs). Benchmarks for this include MultiPL-E^[206], McEval^[207], and VerilogEval^[208]. A popular method to improve performance is to train on manually curated and processing data in low-resource languages such as Coq^[209] and Verilog^[210]. Another line of work aims to achieve transfer between different low-resource languages^{[211][212][213]}. Finally, since the lack of data is a large bottleneck, another popular direction is using relevant

retrievals such as useful functions and library documentation^{[214][215][214]}. For a recent survey of code generation for low-resource languages and DSLs, see^[4].

Security Concerns Surrounding Code Generation: Despite the growing power of LLMs for code generation, their outputs often remain insecure, incorrect, or misaligned with user intent. For instance, BaxBench^[216] evaluates LLMs on generating secure and correct back-ends, revealing that while the average functional correctness is already modest ($\sim 60\%$), the rate of secure outputs is even lower ($< 35\%$). To better understand and quantify these limitations, several benchmarks and evaluation suites have been proposed. SecurityEval^[217], SafeCoder^[218], CodeLMSec^[219], CWEval^[220], and CyberSecEval^{[221][222]} each provide distinct lenses on evaluating vulnerabilities, unsafe API usage, or compliance with common weakness enumerations (CWEs). In response, several approaches introduce human-in-the-loop guardrails, where developers can interactively guide, inspect, or constrain the generation process. Dynex^[223], for instance, supports dynamic, step-wise code synthesis with user feedback, enabling real-time correction and iterative refinement before errors can accumulate.

Human Interaction in Code Generation: Modern code LLMs typically support interactive code generation through conversational interfaces.^[224] conducted a quantitative analysis of developer-ChatGPT interactions using the DevGPT dataset^[225], examining how the quality of the initial prompt influences conversation length. Code LLMs can be further optimized for various interactive scenarios, including debugging environments^[226], educational settings^{[227][228][229][230]}, and use by non-professional programmers^[231]. Beyond human-driven interactions in chat-based setups, more advanced code generation systems such as coding agents can proactively ask clarifying questions^[157] or generate test cases for users to validate^{[232][233]} before generating the actual code, helping to resolve ambiguities.

A.2. Code Transformation

Code Refactoring: Code refactoring aims to simplify and remove repetitions in complex repositories without altering high-level program intent. While there have been traditional methods^[234] that refactor data structures, Aider AI introduces a refactoring benchmark⁶ evaluating LLM's ability to output long chunks of code that simplify complex programs without changing its behavior. More recently, RefactorBench^[52] introduced a more complex benchmark with natural language refactor requests, as well as an LLM agent that can perform refactoring.

Code Migration: Compared to code refactoring, code migration typically refers to mid-scale modifications that affect a program's interface, dependencies, or underlying architecture. Common examples include switching the back-end database from MySQL to PostgreSQL, migrating a machine learning model from TensorFlow to PyTorch, or upgrading the Java version from legacy Java 8 to a more modern Java 17. While recent work has introduced benchmark designed to evaluate library migrations^[235], works at Google^[27] and Amazon^[236] have explored LLM-driven solutions for simple but vast migrations. Google's system identifies locations for changes, generates edits with LLMs fine-tuned on internal code, and automatically validates changes through compilation and test execution.

Code Translation (Transpilation): Moving beyond code migration, transpilation involves large-scale transformation of a program's underlying programming language. Transpilation serves not only to modernize outdated codebases but also to eliminate classes of safety issues inherent to older languages. A particularly active area of research involves transpiling C-based systems to Rust, a systems-level language that offers strong memory and concurrency safety guarantees. This direction has garnered attention, including from the U.S. Department of Defense⁷, which maintains critical infrastructure built on aging C code. An end-to-end LLM-based approach, such as Flourine^[237], has been proposed for real-world code translation, but it has achieved only limited success due to frequent compilation errors. Recent efforts like Syzygy^[33], C2SaferRust^[34], and AlphaTrans^[35] have shown the potential for hybrid approaches combining LLMs with traditional program analysis techniques. However, some significant challenges remain, as identified by^[238], including ensuring correctness in large codebases while maintaining desirable attributes such as speed, reduced vulnerabilities, and idiomaticity. Specifically, We anticipate that the techniques discussed in Section A.3 may help address these remaining challenges.

Code Optimization: Certain refactoring or transpilation tasks are specifically aimed at optimizing code performance. Prior work has explored the use of LLMs for optimizing standalone programs, such as PIE^[119], which targets C++ functions, and AlphaDev^[239], which discovers more efficient sorting algorithms at the assembly level. These tasks are particularly challenging due to the vast search space of possible code transformations. More recently, KernelBench^[37] introduced a benchmark focused on optimizing machine learning models written in high-level PyTorch code into low-level, high-performance CUDA GPU kernels. For a broader overview of language models applied to code optimization, see the survey by^[240].

A.3. Software Testing and Program Analysis

Short-horizon Testing: For short-horizon testing such as unit tests^[241] and property-based tests^[242], LLMs are employed to automatically generate targeted test cases^{[243][244]}, and even hill-climb on code coverage to improve test effectiveness^[245]. At the granularity of individual functions, LLM-generated tests have also been employed to support downstream tasks such as filtering implementations based on behavioral correctness^{[246][247]}, as well as assisting in program debugging by surfacing inputs that expose incorrect behavior^[248].

Long-horizon Testing: Long-horizon testing involves evaluating system behavior across extended executions, complex interactions, or multiple components, potentially embedded within a CI/CD (Continuous Integration or Delivery) pipeline. Fuzzing^[249] is a long-horizon testing approach that continuously generates novel random input. Recent works such as Fuzz4All^[250], KernelGPT^[251], and OSS-Fuzz^{[252][41]} have shown that LLMs can significantly improve effectiveness through better input generation and exploration strategies. Specifically, OSS-Fuzz-Gen^[253] employs diverse LLMs for fuzzing harness generation, helping to find novel and complex crashing interactions.

Static Analysis for Vulnerability Detection: Vulnerability Detection refers to the task of identifying weaknesses or flaws in software code that could be exploited to compromise the system's security, stability, or correctness. A wide range of prior work leverages machine learning models such as Graph Neural Networks (GNNs) and Recurrent Neural Networks (RNNs) to detect software vulnerabilities^{[254][255][256][257][258]}. While some recent methods pre-train or fine-tune LLMs on code-specific datasets^{[259][260][261]} to improve vulnerability classification, several studies have highlighted the limitations of LLMs in real-world software^{[262][263][264]}. To combat such limitations, works like^[265], IRIS^[266], LLMDFE^[267], and InferROI^[268] explored augmenting static analysis tools (e.g., CodeQL) with LLMs for taint and resource leak analyses. More recently,^[71] demonstrated the potential of using LLMs at a much bigger scale by finding a real SQLite vulnerability through exploratory variant analysis.

Specialized Program Analysis: Beyond long-running analysis to identify vulnerabilities, several traditional program analyses have struggled to scale in practice despite their theoretical promise. For instance, inferring program invariants (properties deemed to always be true at a program point) has been challenging with traditional symbolic methods such as Daikon^{[269][270]} while being valuable for exposing bugs^[271] and aiding software evolution^[272]. Similarly, type inference for dynamically typed languages

suffers from coverage limitations of rule-based approaches and requires specialized tools like ShapeIt^[273] for domain-specific challenges such as inferring symbolic tensor shapes.

Specification Inference: Specification inference is the task of automatically recovering formal description of a program’s expected behavior, including pre-conditions, post-conditions, or invariants. The availability of specification is at the core of establishing trust^[274], and existing works^{[275][276]} have shown that LLMs can help the inference of such specifications. For instance,^[277] presents a program structure aware technique for synthesizing pre-conditions for arbitrary code snippets, and have established a dataset of 18K LLM generated pre-conditions on real Java projects.

Invariant Inference: As a subtask of specification inference, invariant inference aims at inferring loop, function, or class invariants, which are greatly helpful in automatic program verification. There have been several LLM-based approaches for invariant identification. They enhance traditional approaches through structured representations^[278], LLM-based prompting^{[279][117]} and re-ranking^[280], and reinforcement learning^[281]. Similarly, works have used sequence-to-sequence models^[282], few-shot LLM approaches like TypeGen^[283], and generate-then-rank methods like TIGER^[284] for type inference. Consequently, we observe new benchmarks emerging in the space such as LIG-MM^[285] for loop-invariant detection.

Binary Analysis: While the aforementioned tasks primarily focus on human-readable programming languages, many can also be extended to operate on compiled machine code, or binaries. One prominent example is binary type inference, which aims to recover high-level type information from low-level binary code. It has seen significant improvements with deep learning models and LLMs^{[286][287]}. These advancements, alongside other LLM-based analyses, have enhanced the capabilities of decompilers, enabling them to synthesize human-readable code from binaries^[288]. Beyond decompilation, LLMs have also been applied to detect security vulnerabilities in binaries^[289] and to generate semantic summaries that capture the high-level intent of binary code^[290].

A.4. Software Maintenance

Code Navigation: Code navigation refers to the task of locating a specific position within a code repository based on either a natural language description^[291] or a programmatic specification^[292]. Common use cases include identifying where a particular functionality is implemented, tracing the origin of user input that leads to a vulnerability, or locating relevant files when starting work on a new

feature. This capability underpins many downstream tasks such as software testing, vulnerability detection, program repair, and code question answering. Code navigation or code search modules are integral components of modern code agents^{[66][293][294]}, often implemented using find commands, embedding-based similarity search, or query-based tools like CodeQL and Semgrep.

Code Documentation and Summarization: Several works have used LLMs for code summarization invoking techniques like prompting^{[187][295][296][297]}. RepoAgent^[298] is a framework that analyzes global contextual relationships in source code to generate fine-grained documentation.^[299] show that LMs are capable of generating good natural language outlines – text descriptions alongside code to partition it into semantically coherent sections. One challenge is that the evaluation of this task is very tricky: the academic community currently lacks datasets and benchmarks that contain good documentation and the automatic evaluation metrics do not align well with human metrics^[300].

Pull Request (PR) Review: In industry, autonomous software agents such as OpenHands^[67] and Devin have been able to automatically review and even fix PRs. At ByteDance, BitsAI-CR^[301] is a code review system that identifies issues based on a manually crafted taxonomy of review rules. In the academic community, there have been several works studying the ability of AI systems to automatically review PRs^{[302][303][304][305]}. Recently, AutoCodeRover^[306] combines LLMs with code search to automatically fix GitHub issues.

Program Repair: Automated program repair has had a long history, with many benchmarks covering different scopes and languages. These include DroixBench^[307] for android apps; Defects4J^[308], GitBug-Java^[309], and growingBugs^{[310][311][312]} for real-world Java; Bugsinpy^[313] for Python; BugSwarm^[314] for multilingual; DebugBench^[315], LiveCodeBench^[23], and Codeflaws^[316] for LeetCode-style problems; and many more.

Historically, there have been many techniques for this task, including heuristic-based APR (using genetic programming to explore the search space of the correct patch), constraint-based APR (treating repair as a constraint-solving task), pattern-based APR (apply expert hand-crafted repair templates), and learning-based APR (using language models)^[317]. More recently, with LLMs, there have been agent-based approaches such as FixAgent^[318] using agents specializing in different aspects of debugging, and RepairAgent^[293] that invokes suitable tools. On the other hand, Agentless^[294] uses a three-phase process of localization, repair, and patch validation.

Finally, program repair has also been used as a tool to improve code generation, where error messages and incorrect test cases are fed back into the model to improve code generation^{[319][320][321][322][323][324]}. This is also known as self-repair or self-debugging. For a much more comprehensive survey of automated program repair, we recommend the reader check out [this website](#)⁸.

Code Understanding and Question Answering: Code understanding with language models has been studied for many years. In earlier days, researchers used the CodeXGLUE^[325] benchmark containing tasks such as clone detection, code search, code summarization, and so on.^[326] create an IDE plugin containing features that help users understand code through explaining highlighted sections of code and explaining domain-specific code.^[327] present a survey touching on reasoning-enhanced code intelligence.

A.5. Scaffolding and Meta-Code

Beyond code generation, the broader software engineering ecosystem includes DevOps workflows, CI/CD pipelines, and Infrastructure-as-Code (IaC). LLMs have shown particular promise in generating, debugging, and explaining CI/CD configurations (e.g., GitHub Actions, Jenkinsfiles), assisting with environment setup, test orchestration, and deployment logic. A case study at Ericsson^[328] demonstrates how an LLM-based chatbot can support CI/CD question answering, enabling engineers to better understand and manage deployment pipelines. LLMs are also being explored for automated testing across heterogeneous software environments. ExecutionAgent^[329] presents a language model-driven agent that autonomously installs, configures, and runs test suites for arbitrary projects.

Beyond CI/CD and testing, LLMs are increasingly used to reason about configuration logic and scaffolding code, which is a critical but often overlooked layer of modern software systems. For instance, ^[330] conducted an empirical study of real-world configuration errors, identifying systemic causes of failure such as external dependencies, inter-parameter violations, and overlooked default parameters. Building on this line of work, Ciri^[44] confirms the feasibility of using LLMs for configuration validation. Further, in the domain of IaC, an empirical study of 812 open-source Terraform projects found that while access policies are commonly adopted, critical practices like encryption at rest are often neglected^[331]. This highlights the opportunity for LLMs to assist practitioners in detecting and enforcing security best practices in IaC configurations.

A.6. Formal Verification

There are a variety of programming languages designed with different principles to support formal verification. Some of the popular ones include TLA^[332], Coq^[48], Lean^[110], Dafny^[333], Isabelle^[334], and Verus^[335].

Formal software verification has seen a few great successes in the last few years: Astrée^[336] was able to completely verify that Airbus A340's primary flight-control software had no run-time errors, verifying 132,000 lines of C code. More recently, formal methods have been applied to verify a cryptographic server^[337] and an IoT lightbulb at both a hardware and software level^[338]. CompCert^[47], a verified compiler and seL4^[339], a verified microkernel are demonstrations that formal methods could be promising for verifiable code. At Amazon, formal methods been used to verify and protect cryptographic software^[340], cloud resources^[341], and authorization^[342]. Notably, SV-COMP^[343] is an annual competition designed to evaluate program verifiers using a curated benchmark of verifiable C and Java code. It even includes samples from the Linux Driver Verification (LDV) project^[344], aiding the verification of Linux kernel device drivers. For more applications, we refer the reader to the survey in^[345].

Recently, the ability of LLMs to write formal verification code. Benchmarks like DafnyBench^[346] and miniCodeProps^[347] were designed to measure the ability of LLMs to write software proofs in Dafny and Lean, respectively. In Dafny,^[348] use a combination of search and prompting to create a synthetic dataset of annotations greatly improving performance on DafnyBench. Clover^[56] generates code alongside consistency checks (like Dafny annotations),^[349] employ Dafny as an intermediate language to improve code generation, and^[350] explore prompting and retrieval to generate Dafny. In Rust, Verus is a popular formal verification language, with AutoVerus^[351] and AlphaVerus^[352] generating verified specifications and proofs for Rust functions. There are also many IDE plugins designed to help humans to write code in formal languages such as Dafny and Lean such as^[353], Lean Copilot^[354], and llmstep^[355].

Finally, there is a growing interest of work in using formal languages like Lean for mathematical theorem proving, which is covered comprehensively in^[356] and^[6].

Acknowledgements

We thank Alex Polozov, Baptiste Roziere, Daya Guo, Jenny Liang, Jiawei Liu, Justin Chiu, Kexun Zhang, Leonardo Hernandez Cano, Li Zhong, Michael Wang, Silas Alberti, Theo Olausson, Valerie Chen, Xingyao

Wang, Yangruibo Ding, Yuxiang Wei, Zhiruo Wang, and several anonymous workshop reviewers for providing valuable feedback regarding various stages of the draft.

We also thank the following people for bringing up illustrative examples mentioned in this paper: Silas Alberti (debugging cloud applications), Chuyue Sun (incomplete specification in Verus), MIT's 6.172 Course (performance instrumentation), Theo Olausson (costly disasters), Songlin Yang (syntax error in Triton).

A. Gu is supported by the National Science Foundation (NSF) Graduate Research Fellowship under Grant No. 2141064. N. Jain is supported by NSF grants CCF:1900968, CCF:1908870, and by SKY Lab industrial sponsors and affiliates. A. Solar-Lezama is supported by the National Science Foundation (NSF) and Intel Corporation through NSF Grant CCF:2217064. D. Yang is supported by the ONR YIP Award N000142412532.

Footnotes

¹ We follow page 9, Table 2 from their paper

² <https://www.kprize.ai/>

³ As reported by the [BigCode Models Leaderboard](#) on the MultiPL-E benchmark^[206]

⁴ <https://www.cursor.so>

⁵ <https://www.tabnine.com>

⁶ <https://github.com/Aider-AI/refactor-benchmark>

⁷ <https://www.darpa.mil/news/2024/memory-safety-vulnerabilities>

⁸ <https://program-repair.org/>

References

- ¹ [Liang JT, Yang C, Myers BA \(2024\)](#). "A large-scale survey on the usability of ai programming assistants: Successes and challenges." In: *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering*. pp. 1–13.
- ² [Sergeyuk A, Golubev Y, Bryksin T, Ahmed I \(2025\)](#). "Using AI-based coding assistants in practice: State of affairs, perceptions, and ways forward". *Information and Software Technology*. **178**: 107610.

3. ^aWang J, Huang Y, Chen C, Liu Z, Wang S, Wang Q. "Software testing with large language models: Survey, landscape, and vision". *IEEE Transactions on Software Engineering*. 2024.
4. ^a, ^bJoel S, Wu JJ, Fard FH (2024). "A Survey on LLM-based Code Generation for Low-Resource and Domain-Specific Programming Languages". *arXiv preprint arXiv:2410.03981*. Available from: <https://arxiv.org/abs/2410.03981>.
5. ^aZhang Q, Fang C, Ma Y, Sun W, Chen Z (2023). "A survey of learning-based automated program repair". *ACM Transactions on Software Engineering and Methodology*. 33 (2): 1–69.
6. ^a, ^bYang K, Poesia G, He J, Li W, Lauter K, Chaudhuri S, Song D (2024). "Formal Mathematical Reasoning: A New Frontier in AI". *arXiv preprint arXiv:2412.16075*. Available from: <https://arxiv.org/abs/2412.16075>.
7. ^aFan A, Gokkaya B, Harman M, Lyubarskiy M, Sengupta S, Yoo S, Zhang JM. Large language models for software engineering: Survey and open problems. In: *2023 IEEE/ACM International Conference on Software Engineering: Future of Software Engineering (ICSE-FoSE)*. IEEE; 2023. p. 31-53.
8. ^aOzkaya I. Application of large language models to software engineering tasks: Opportunities, risks, and implications. *IEEE Software*. 40(3):4–8, 2023.
9. ^aWong MF, Guo S, Hang CN, Ho SW, Tan CW (2023). "Natural language generation and understanding of big code for AI-assisted programming: A review". *Entropy*. 25 (6): 888.
10. ^aZheng Z, Ning K, Wang Y, Zhang J, Zheng D, Ye M, Chen J (2023). "A survey of large language models for code: Evolution, benchmarking, and future trends". *arXiv preprint arXiv:2311.10372*. Available from: <https://arxiv.org/abs/2311.10372>.
11. ^aHou X, Zhao Y, Liu Y, Yang Z, Wang K, Li L, Luo X, Lo D, Grundy J, Wang H (2024). "Large language models for software engineering: A systematic literature review". *ACM Transactions on Software Engineering and Methodology*. 33 (8): 1–79.
12. ^aJin H, Huang L, Cai H, Yan J, Li B, Chen H (2024). "From llms to llm-based agents for software engineering: A survey of current, challenges and future". *arXiv preprint arXiv:2408.02479*. [arXiv:2408.02479](https://arxiv.org/abs/2408.02479).
13. ^aWan Y, Bi Z, He Y, Zhang J, Zhang H, Sui Y, Xu G, Jin H, Yu P (2024). "Deep Learning for Code Intelligence: Survey, Benchmark and Toolkit". *ACM Computing Surveys*. ACM New York, NY.
14. ^aRoychoudhury A, Pasareanu C, Pradel M, Ray B (2025). "AI Software Engineer: Programming with Trust". *arXiv preprint arXiv:2502.13767*.
15. ^a, ^b, ^c, ^dChen M, Tworek J, Jun H, Yuan Q, Pinto HP de Oliveira, Kaplan J, Edwards H, Burda Y, Joseph N, Brockman G, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*. 2021.

16. ^a ^b ^c Austin J, Odena A, Nye M, Bosma M, Michalewski H, Dohan D, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*. 2021.
17. ^Δ Liu S, Zhu H, Liu J, Xin S, Li A, Long R, Chen L, Yang J, Xia J, Peng ZY, Liu S, Zhang Z, Zhang G, Huang W, Shen K, Xiang L (2024). "FullStack Bench: Evaluating LLMs as Full Stack Coders". *arXiv*. Available from: <https://arxiv.org/abs/2412.00535>.
18. ^Δ Zhuo TY, Vu MC, Chim J, Hu H, Yu W, Widyasari R, Yusuf INB, Zhan H, He J, Paul I, et al. Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions. *arXiv preprint arXiv:2406.15877*. 2024.
19. ^a ^b Zhao W, Jiang N, Lee C, Chiu JT, Cardie C, Gall{e} M, Rush AM (2024). "Commit0: Library Generation from Scratch". *arXiv preprint arXiv:2412.01769*.
20. ^a ^b ^c Jimenez CE, Yang J, Wettig A, Yao S, Pei K, Press O, Narasimhan KR. "SWE-bench: Can language models resolve real-world github issues?" In: *The Twelfth International Conference on Learning Representations*; 2024. Available from: <https://openreview.net/forum?id=VTF8yNQm66>.
21. ^Δ Hendrycks D, Basart S, Kadavath S, Mazeika M, Arora A, Guo E, Burns C, Puranik S, He H, Song D, Steinhardt J (2021). "Measuring coding challenge competence with apps". *NeurIPS*. 2021.
22. ^Δ Li Y, Choi D, Chung J, Kushman N, Schrittwieser J, Leblond R, Eccles T, Keeling J, Gimeno F, Dal Lago A, et al. Competition-level code generation with alphacode. *Science*. **378** (6624): 1092–1097, 2022.
23. ^a ^b ^c Jain N, Han K, Gu A, Li WD, Yan F, Zhang T, Wang S, Solar-Lezama A, Sen K, Stoica I (2024). "LiveCodeBench: Holistic and Contamination Free Evaluation of Large Language Models for Code". *arXiv*. [arXiv:2403.07974](https://arxiv.org/abs/2403.07974).
24. ^a ^b Treude C, Gerosa MA (2025). "How Developers Interact with AI: A Taxonomy of Human-AI Collaboration in Software Engineering". *arXiv preprint arXiv:2501.08774*.
25. ^Δ Morris MR, Sohl-Dickstein J, Fiedel N, Warkentin T, Dafoe A, Faust A, Farabet C, Legg S (2023). "Levels of AGI: Operationalizing Progress on the Path to AGI". *arXiv preprint arXiv:2311.02462*. Available from: <https://arxiv.org/abs/2311.02462>.
26. ^Δ Parnas DL (1972). "On the criteria to be used in decomposing systems into modules". *Communications of the ACM*. **15** (12): 1053–1058.
27. ^a ^b Nikolov S, Codecasa D, Sjovald A, Tabachnyk M, Chandra S, Taneja S, Ziftci C (2025). "How is Google using AI for internal code migrations?" *arXiv preprint arXiv:2501.06972*.
28. ^Δ Pierre Ricadat. Scala 3 migration: Report from the field. <https://blog.pierre-ricadat.com/scala-3-migration-report-from-the-field>, 2025.

29. [△]Taulli T (2020 Jul). "COBOL LANGUAGE: Call it a comeback?" *Forbes*. Available from: <https://www.forbes.com/sites/tomtaulli/2020/07/13/cobol-language-call-it-a-comeback/>.
30. [△]Sneed HM. "Extracting business logic from existing COBOL programs as a basis for redevelopment." In: *Proceedings 9th International Workshop on Program Comprehension. IWPC 2001. IEEE*; 2001. p. 167–175.
31. [△]Sellink A, Sneed H, Verhoef C (2002). "Restructuring of COBOL/CICS legacy systems". *Science of Computer Programming*. **45** (2–3): 193–243.
32. [△]Sneed HM. "Migrating from COBOL to Java." In: *2010 IEEE International Conference on Software Maintenance*. IEEE; 2010. p. 1–7.
33. [△][♢]Shetty M, Jain N, Godbole A, Seshia SA, Sen K (2024). "Syzygy: Dual Code-Test C to (safe) Rust Translation using LLMs and Dynamic Analysis". *arXiv preprint arXiv:2412.14234*.
34. [△][♢]Nitin V, Krishna R, Valle LL, Ray B (2025). "C2SaferRust: Transforming C Projects into Safer Rust with NeuroSymbolic Techniques". *arXiv preprint arXiv:2501.14257*.
35. [△][♢]Ibrahimzada AR, Ke K, Pawagi M, Abid MS, Pan R, Sinha S, Jabbarvand R (2024). "Repository-level compositional code translation and validation". *arXiv preprint arXiv:2410.24117*.
36. [△]Zhao C, Zhou S, Zhang L, Deng C, Xu Z, Liu Y, Yu K, Li J, Zhao L (2025). "DeepEP: an efficient expert-parallel communication library". <https://github.com/deepseek-ai/DeepEP>.
37. [△][♢][♣][♠]Ouyang A, Guo S, Mirhoseini A (2024). "KernelBench: Can LLMs Write GPU Kernels?". <https://scalingintelligence.stanford.edu/blogs/kernelbench/>.
38. [△]Chromium (2018). "10 years of speed in chrome". *Chromium Blog*. Available from: <https://blog.chromium.org/2018/09/10-years-of-speed-in-chrome-11.html>.
39. [△]Mosolygó B, Vándor N, Antal G, Hegedüs P (2021). "On the rise and fall of simple stupid bugs: a life-cycle analysis of sstubs". In: *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. IEEE. pp. 495–499.
40. [△]Trent B, Li A (2025). "Concurrency bugs in Lucene: How to fix optimistic concurrency failures". <https://www.elastic.co/search-labs/blog/optimistic-concurrency-lucene-debugging>.
41. [△][♢]Chang O, Liu D, Metzman J, Google Open Source Security Team. "Leveling Up Fuzzing: Finding more vulnerabilities with AI". <https://security.googleblog.com/2024/11/leveling-up-fuzzing-finding-more.html>, 2024.
42. [△]Foster C, Gulati A, Harman M, Harper I, Mao K, Ritchey J, Robert H, Sengupta S (2025). "Mutation-Guided LLM-based Test Generation at Meta". *arXiv preprint arXiv:2501.12862*.
43. [△]Hawkes B (2019). "Oday 'In the Wild'". <https://googleprojectzero.blogspot.com/p/oday.html>.

44. ^a ^bLian X, Chen Y, Cheng R, Huang J, Thakkar P, Zhang M, Xu T. "Large Language Models as Configuration Validators." In: 2025 IEEE/ACM 47th International Conference on Software Engineering (ICSE). IEEE Computer Society; 2024. p. 204–216.
45. [△]Terrateam (2024). "Using LLMs to Generate Terraform Code". <https://terrateam.io/blog/using-llms-to-generate-terraform-code/>.
46. [△]Zee K, Kuncak V, Rinard M (2008). "Full functional verification of linked data structures". ACM SIGPLAN Notices. 43 (6): 349–361.
47. ^a ^bLeroy X, Blazy S, Kästner D, Schommer B, Pister M, Ferdinand C (2016). "CompCert—a formally verified optimizing compiler". In: ERTS 2016: Embedded Real Time Software and Systems, 8th European Congress.
48. ^a ^bThe Coq Development Team. The Coq Reference Manual -- Release 8.19.0. 2024. Available from: <https://coq.inria.fr/doc/V8.19.0/refman>.
49. [△]Bessey A, Block K, Chelf B, Chou A, Fulton B, Hallem S, Henri-Gros C, Kamsky A, McPeak S, Engler D (2010). "A few billion lines of code later: using static analysis to find bugs in the real world". Communications of the ACM. 53 (2): 66–75.
50. ^a ^bChi W, Chen V, Angelopoulos AN, Chiang W-L, Mittal A, Jain N, Zhang T, Stoica I, Donahue C, Talwalkar A (2025). "Copilot Arena: A Platform for Code LLM Evaluation in the Wild". arXiv. Available from: <https://arxiv.org/abs/2502.09328>.
51. [△]Jain K, Synnaeve G, Rozière B (2024). "Testgeneval: A real world unit test generation and test completion benchmark". arXiv preprint arXiv:2410.00752.
52. ^a ^bGautam D, Garg S, Jang J, Sundaresan N, Moghaddam RZ (2024). "RefactorBench: Evaluating Stateful Reasoning In Language Agents Through Code". In: NeurIPS 2024 Workshop on Open-World Agents.
53. [△]Miserendino S, Wang M, Patwardhan T, Heidecke J (2025). "SWE-Lancer: Can Frontier LLMs Earn \$1 Million from Real-World Freelance Software Engineering?" arXiv preprint arXiv:2502.12115.
54. [△]Liu J, Xie S, Wang J, Wei Y, Ding Y, Zhang L. Evaluating language models for efficient code generation. In: First Conference on Language Modeling; 2024. Available from: <https://openreview.net/forum?id=IBCBMeAhmC>.
55. [△]Mei X, Singaria PS, Del Castillo J, Xi H, Bao T, Wang R, Shoshitaishvili Y, Doupé A, Pearce H, Dolan-Gavitt B, et al. ARVO: Atlas of Reproducible Vulnerabilities for Open Source Software. arXiv preprint arXiv:2408.02153. 2024.
56. ^a ^bSun C, Sheng Y, Padon O, Barrett C. "Clover: Closed-Loop Verifiable Code Generation." In: International Symposium on AI Verification. Springer; 2024. p. 134-155.

57. ^a Xu C, Guan S, Greene D, Kechadi M, et al. Benchmark data contamination of large language models: A survey. *arXiv preprint arXiv:2406.04244*. 2024.
58. ^Δ Aleithan R, Xue H, Mohajer MM, Nnorom E, Uddin G, Wang S (2024). "Swe-bench+: Enhanced coding benchmark for llms". *arXiv preprint arXiv:2410.06992*.
59. ^Δ Matton A, Sherborne T, Aumiller D, Tommasone E, Alizadeh M, He J, Ma R, Voisin M, Gilsenan-McMahon E, Gallé M (2024). "On leakage of code generation evaluation datasets". *arXiv preprint arXiv:2407.07565*.
60. ^Δ Riddell M, Ni A, Cohan A (2024). "Quantifying contamination in evaluating code generation capabilities of language models". *arXiv preprint arXiv:2403.04811*. Available from: <https://arxiv.org/abs/2403.04811>.
61. ^Δ Schick T, Dwivedi-Yu J, Dessì R, Raileanu R, Lomeli M, Hambro E, Zettlemoyer L, Cancedda N, Scialom T (2023). "Toolformer: Language models can teach themselves to use tools". *Advances in Neural Information Processing Systems*. **36**: 68539–68551.
62. ^Δ Patil SG, Zhang T, Wang X, Gonzalez JE (2023). "Gorilla: Large Language Model Connected with Massive APIs". *arXiv preprint arXiv:2305.15334*. Available from: [arXiv:2305.15334](https://arxiv.org/abs/2305.15334).
63. ^Δ Olausson TX, Inala JP, Wang C, Gao J, Solar-Lezama A (2023). "Is self-repair a silver bullet for code generation?" In: *The Twelfth International Conference on Learning Representations*.
64. ^Δ Zhong L, Wang Z, Shang J (2024). "Ldb: A large language model debugger via verifying runtime execution step-by-step". *arXiv preprint arXiv:2402.16906*.
65. ^Δ Gehring J, Zheng K, Copet J, Mella V, Cohen T, Synnaeve G (2024). "Rlef: Grounding code llms in execution feedback with reinforcement learning". *arXiv preprint arXiv:2410.02089*.
66. ^a Yang J, Jimenez CE, Wettig A, Lieret K, Yao S, Narasimhan KR, Press O (2024). "SWE-agent: Agent-Computer Interfaces Enable Automated Software Engineering". In: *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. Available from: <https://arxiv.org/abs/2405.15793>.
67. ^a Wang X, Li B, Song Y, Xu FF, Tang X, Zhuge M, Pan J, Song Y, Li B, Singh J, Tran HH, Li F, Ma R, Zheng M, Qian B, Shao Y, Muennighoff N, Zhang Y, Hui B, Lin J, Brennan R, Peng H, Ji H, Neubig G (2024). "OpenHands: An Open Platform for AI Software Developers as Generalist Agents". *arXiv*. Available from: <https://arxiv.org/abs/2407.16741>.
68. ^Δ Anthropic (2024). "Raising the bar on SWE-bench Verified with Claude 3.5 Sonnet".
69. ^Δ Liu Y, Gao P, Wang X, Liu J, Shi Y, Zhang Z, Peng C (2024). "MarsCode Agent: AI-native Automated Bug Fixing". *arXiv preprint arXiv:2409.00899*. Available from: <https://arxiv.org/abs/2409.00899>.
70. ^Δ Bairi R, Sonwane A, Kanade A, Iyer A, Parthasarathy S, Rajamani S, Ashok B, Shet S (2024). "Codeplan: Repository-level coding using llms and planning". *Proceedings of the ACM on Software Engineering*. **1** (FSE): 67

71. ^aGoogle BigSleep (2024). "From Naptime to Big Sleep: Using Large Language Models To Catch Vulnerabilities In Real-World Code". <https://googleprojectzero.blogspot.com/2024/10/from-naptime-to-big-sleep.html>.
72. ^ΔSchardl TB, Denniston T, Doucet D, Kuszmaul BC, Lee IA, Leiserson CE (2017). "The CSI framework for compiler-inserted program instrumentation". *Proceedings of the ACM on Measurement and Analysis of Computing Systems*. 1 (2): 1–25.
73. ^ΔZou W, Xuan J, Xie X, Chen Z, Xu B (2019). "How does code style inconsistency affect pull request integration? An exploratory study on 117 GitHub projects". *Empirical Software Engineering*. 24: 3871–3903.
74. ^ΔWang Y, Jiang T, Liu M, Chen J, Zheng Z (2024). "Beyond functional correctness: Investigating coding style inconsistencies in large language models". *arXiv preprint arXiv:2407.00456*. Available from: <https://arxiv.org/abs/2407.00456>.
75. ^ΔWeisz JD, Kumar S, Muller M, Browne KE, Goldberg A, Heintze E, Bajpai S (2024). "Examining the Use and Impact of an AI Code Assistant on Developer Productivity and Experience in the Enterprise". *arXiv preprint arXiv:2412.06603*.
76. ^ΔShao Y, Samuel V, Jiang Y, Yang J, Yang D (2024). "Collaborative Gym: A Framework for Enabling and Evaluating Human-Agent Collaboration". *arXiv preprint arXiv:2412.15701*. [arXiv:2412.15701](https://arxiv.org/abs/2412.15701).
77. ^ΔNahar N, Zhou S, Lewis G, Ke4stner C. "Collaboration challenges in building ml-enabled systems: Communication, documentation, engineering, and process." In: *Proceedings of the 44th international conference on software engineering*. 2022. p. 413–425.
78. ^ΔWang R, Cheng R, Ford D, Zimmermann T (2024). "Investigating and Designing for Trust in AI-powered Code Generation Tools". In: *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency, FAccT '24, Rio de Janeiro, Brazil*. New York, NY, USA: Association for Computing Machinery. p. 1475–1493. doi:[10.1145/3630106.3658984](https://doi.org/10.1145/3630106.3658984).
79. ^ΔBenderius O, Berger C, Lundgren VM (2017). "The best rated human-machine interface design for autonomous vehicles in the 2016 grand cooperative driving challenge". *IEEE Transactions on intelligent transportation systems*. 19 (4): 1302–1307.
80. ^ΔTinga AM, Cleij D, Jansen RJ, van der Kint S, van Nes N (2022). "Human machine interface design for continuous support of mode awareness during automated driving: An online simulation". *Transportation research part F: traffic psychology and behaviour*. 87: 102–119.
81. ^a^bEllis K, Wong C, Nye M, Sablé-Meyer M, Morales L, Hewitt L, Cary L, Solar-Lezama A, Tenenbaum JB. "Dreamcoder: Bootstrapping inductive program synthesis with wake-sleep library learning." In: *Proceedings of*

the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation. 2021. p. 835–850.

82. [△]Stengel-Eskin E, Prasad A, Bansal M. REGAL: refactoring programs to discover generalizable abstractions. In: *Proceedings of the 41st International Conference on Machine Learning, ICML'24*. JMLR.org; 2024.
83. [△]Bowers M, Olausson TX, Wong L, Grand G, Tenenbaum JB, Ellis K, Solar-Lezama A. Top-down synthesis for library learning. *Proc. ACM Program. Lang.*. 2023 Jan; 7(POPL):41. doi:[10.1145/3571234](https://doi.org/10.1145/3571234).
84. [△]Kraska T, Beutel A, Chi EH, Dean J, Polyzotis N (2018). "The case for learned index structures". In: *Proceedings of the 2018 international conference on management of data*. pp. 489–504.
85. [△]Hawkins P, Aiken A, Fisher K, Rinard M, Sagiv M (2011). "Data representation synthesis". In: *Proceedings of the 32nd ACM SIGPLAN conference on Programming language design and implementation*. pp. 38–49.
86. [△][‡]Jain N, Shetty M, Zhang T, Han K, Sen K, Stoica I (2024). "R2e: Turning any github repository into a programming agent environment". *Forty-first International Conference on Machine Learning*.
87. [△]Berlot-Attwell I, Rudzicz F, Si X (2024). "Library Learning Doesn't: The Curious Case of the Single-Use Library". *arXiv preprint arXiv:2410.20274*.
88. [△]Potvin R, Levenberg J (2016). "Why Google stores billions of lines of code in a single repository". *Communications of the ACM*. 59 (7): 78–87.
89. [△]Gao Y, Xiong Y, Gao X, Jia K, Pan J, Bi Y, Dai Y, Sun J, Wang H (2023). "Retrieval-augmented generation for large language models: A survey". *arXiv preprint arXiv:2312.10997*. [arXiv:2312.10997](https://arxiv.org/abs/2312.10997).
90. [△]Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, Küttler H, Lewis M, Yih W, Rocktäschel T, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*. 33:9459–9474, 2020.
91. [△]Ma W, Liu S, Zhao M, Xie X, Wang W, Hu Q, Zhang J, Liu Y (2024). "Unveiling code pre-trained models: Investigating syntax and semantics capacities". *ACM Transactions on Software Engineering and Methodology*. 33 (7): 1–29.
92. [△][‡]Utpala S, Gu A, Chen PY (2023). "Language Agnostic Code Embeddings". *arXiv preprint arXiv:2310.16803*. Available from: <https://arxiv.org/abs/2310.16803>.
93. [△]Wang ZZ, Asai A, Yu XV, Xu FF, Xie Y, Neubig G, Fried D (2024). "Coderag-bench: Can retrieval augment code generation?" *arXiv preprint arXiv:2406.14497*.
94. [△]Su H, Yen H, Xia M, Shi W, Muennighoff N, Wang H, Liu H, Shi Q, Siegel ZS, Tang M, et al. Bright: A realistic and challenging benchmark for reasoning-intensive retrieval. *arXiv preprint arXiv:2407.12883*. 2024.

95. [△]Yang J, Jimenez CE, Zhang AL, Lieret K, Yang J, Wu X, Press O, Muennighoff N, Synnaeve G, Narasimhan KR, et al. SWE-bench Multimodal: Do AI Systems Generalize to Visual Software Domains? arXiv preprint arXiv:2410.03859. 2024.
96. [△]Ding Y, Wang Z, Ahmad W, Ding H, Tan M, Jain N, Ramanathan MK, Nallapati R, Bhatia P, Roth D, et al. (2023). "Crosscodeeval: A diverse and multilingual benchmark for cross-file code completion". *Advances in Neural Information Processing Systems*. 36: 46701–46723.
97. [△]Roychoudhury A, Zeller A (2025). "Will AI replace Software Engineers? Hold your Breath". arXiv preprint arXiv:2502.20429.
98. [△]Gu A, Rozière B, Leather H, Solar-Lezama A, Synnaeve G, Wang SI (2024). "CRUXEval: A Benchmark for Code Reasoning, Understanding and Execution". arXiv preprint arXiv:2401.03065. 2024. Available from: <http://arxiv.org/abs/2401.03065>.
99. [△]Ni A, Allamanis M, Cohan A, Deng Y, Shi K, Sutton C, Yin P (2024). "Next: Teaching large language model to reason about code execution". arXiv preprint arXiv:2404.14662.
100. [△]Ding Y, Steenhoek B, Pei K, Kaiser G, Le W, Ray B (2024). "Traced: Execution-aware pre-training for source code". In: *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering*. pp. 1–12.
101. [△]Ahmed T, Bird C, Devanbu P, Chakraborty S (2024). "Studying LLM Performance on Closed-and Open-source Data". arXiv preprint arXiv:2402.15100.
102. [△]Blinn A, Li X, Kim JH, Omar C (2024). "Statically contextualizing large language models with typed holes". *Proceedings of the ACM on Programming Languages*. 8 (OOPSLA2): 468–498.
103. [△]Hui B, Yang J, Cui Z, Yang J, Liu D, Zhang L, Liu T, Zhang J, Yu B, Lu K, et al. Qwen2. 5-coder technical report. arXiv preprint arXiv:2409.12186. 2024.
104. [△]Wu T, Luo L, Li YF, Pan S, Vu TT, Haffari G (2024). "Continual learning for large language models: A survey". arXiv preprint arXiv:2402.01364.
105. [△]Wang L, Zhang X, Su H, Zhu J. "A comprehensive survey of continual learning: theory, method and application". *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2024.
106. [△]Liu ZL, Pandit S, Ye X, Choi E, Durrett G (2024). "Codeupdatearena: Benchmarking knowledge editing on api updates". arXiv preprint arXiv:2407.06249.
107. [△]Islah N, Gehring J, Misra D, Muller E, Rish I, Zhuo TY, Caccia M (2024). "GitChameleon: Unmasking the Version-Switching Capabilities of Code Generation Models". arXiv preprint arXiv:2411.05830.
108. [△]Kumarappan A, Tiwari M, Song P, George RJ, Xiao C, Anandkumar A (2024). "LeanAgent: Lifelong Learning for Formal Theorem Proving". arXiv preprint arXiv:2410.06209.

109. [△]Kharm M, Choi S, AlKhanafseh M, Mohaisen D (2025). "Security and Quality in LLM-Generated Code: A Multi-Language, Multi-Model Analysis". *arXiv preprint arXiv:2502.01853*.
110. [△][♢]De Moura L, Kong S, Avigad J, Van Doorn F, von Raumer J. "The Lean theorem prover (system description)". In: *Automated Deduction-CADE-25: 25th International Conference on Automated Deduction*, Berlin, Germany, August 1-7, 2015, *Proceedings 25*. Springer; 2015. p. 378-388.
111. [△]Moura L de, Ullrich S. The Lean 4 theorem prover and programming language. In: *Automated Deduction-CADE 28: 28th International Conference on Automated Deduction*, Virtual Event, July 12--15, 2021, *Proceedings 28*. Springer; 2021. p. 625-635.
112. [△][♢]Mankowitz DJ, Michi A, Zhernov A, Gelmi M, Selvi M, Paduraru C, Leurent E, Iqbal S, Lespiau JB, Ahern A, et al. Faster sorting algorithms discovered using deep reinforcement learning. *Nature*. **618** (7964): 257-263, 2023.
113. [△]Neri C (2023). "Shorter and faster than Sort3AlphaDev". *arXiv preprint arXiv:2307.14503*.
114. [△]Ullrich M, Hack S. Synthesis of sorting kernels. In: *Proceedings of the 23rd ACM/IEEE International Symposium on Code Generation and Optimization, CGO '25*, New York, NY, USA: Association for Computing Machinery; 2025. p. 1-14. ISBN 9798400712753. doi:[10.1145/3696443.3708954](https://doi.org/10.1145/3696443.3708954).
115. [△]Chen H, Ziegler D, Chajed T, Chlipala A, Kaashoek MF, Zeldovich N (2015). "Using Crash Hoare logic for certifying the FSCQ file system". In: *Proceedings of the 25th Symposium on Operating Systems Principles*. 2015: 18-37.
116. [△]Ding Y, Peng J, Min MJ, Kaiser G, Yang J, Ray B (2024). "Semcoder: Training code language models with comprehensive semantics". *arXiv preprint arXiv:2406.01006*. Available from: <https://arxiv.org/abs/2406.01006>.
117. [△][♢]Pei K, Bieber D, Shi K, Sutton C, Yin P. "Can large language models reason about program invariants?" In: *International Conference on Machine Learning*. PMLR; 2023. p. 27496-27520.
118. [△]Guo D, Ren S, Lu S, Feng Z, Tang D, Liu S, Zhou L, Duan N, Svyatkovskiy A, Fu S, et al. (2020). "Graphcodebert: Pre-training code representations with data flow". *arXiv preprint arXiv:2009.08366*. Available from: <https://arxiv.org/abs/2009.08366>.
119. [△][♢]Shypula A, Madaan A, Zeng Y, Alon U, Gardner J, Hashemi M, Neubig G, Ranganathan P, Bastani O, Yazdambakhsh A (2023). "Learning performance-improving code edits". *arXiv preprint arXiv:2302.07867*.
120. [△]Liu J, Zhang L (2025). "Code-R1: Reproducing R1 for Code with Reliable Rewards". Available from: <https://github.com/ganler/code-r1>.
121. [△]Li R, Fu J, Zhang BW, Huang T, Sun Z, Lyu C, Liu G, Jin Z, Li G (2023). "Taco: Topics in algorithmic code generation dataset". *arXiv preprint arXiv:2312.14852*. Available from: <https://arxiv.org/abs/2312.14852>.

122. [△]Gulwani S, Polozov O, Singh R, et al. (2017). "Program synthesis". *Foundations and Trends® in Programming Languages*. 4 (1-2): 1–119.
123. [△]Li WD, Hu K, Larsen C, Wu Y, Alford S, Woo C, Dunn SM, Tang H, Naim M, Nguyen D, et al. (2024). "Combining induction and transduction for abstract reasoning". *arXiv preprint arXiv:2411.02272*.
124. [△]Trinh TH, Wu Y, Le QV, He H, Luong T (2024). "Solving olympiad geometry without human demonstrations". *Nature*. 625 (7995): 476–482.
125. [△]Google (2024). "AI achieves silver-medal standard solving International Mathematical Olympiad problems". Available from: <https://deepmind.google/discover/blog/ai-solves-imo-problems-at-silver-medal-level/>.
126. [△]Chervonyi Y, Trinh TH, Ol{\v{s}}{\a}k M, Yang X, Nguyen H, Menegali M, Jung J, Verma V, Le QV, Luong T (2025). "Gold-medalist Performance in Solving Olympiad Geometry with AlphaGeometry2". *arXiv preprint arXiv:2502.03544*. Available from: <https://arxiv.org/abs/2502.03544>.
127. [△]Kocetkov D, Li R, Ben Allal L, Li J, Mou C, Muñoz Ferrandis C, Jernite Y, Mitchell M, Hughes S, Wolf T, et al. T he stack: 3 tb of permissively licensed source code. *arXiv preprint arXiv:2211.15533*. 2022.
128. [△]Lozhkov A, Li R, Ben Allal L, Cassano F, Lamy-Poirier J, Tazi N, Tang A, Pykhtar D, Liu J, Wei Y, et al. (2024). "Starcoder 2 and the stack v2: The next generation". *arXiv preprint arXiv:2402.19173*.
129. [△]Chandra S (2024). "AI in Software Engineering at Google: Progress and the Path Ahead (Invited Talk)". In: *Proceedings of the 1st ACM International Conference on AI-Powered Software*. pp. 182–182.
130. [△]Murali V, Maddila C, Ahmad I, Bolin M, Cheng D, Ghorbani N, Fernandez R, Nagappan N, Rigby PC (2024). "AI-assisted Code Authoring at Scale: Fine-tuning, deploying, and mixed methods evaluation". *Proceedings of the ACM on Software Engineering*. 1 (FSE): 1066–1085.
131. [△]Ziegler A, Kalliamvakou E, Li XA, Rice A, Rifkin D, Simister S, Sittampalam G, Aftandilian E (2024). "Measuring GitHub Copilot's impact on productivity". *Communications of the ACM*. 67 (3): 54–63.
132. [△]Li J, Guo D, Yang D, Xu R, Wu Y, He J (2025). "CodeI/O: Condensing Reasoning Patterns via Code Input-Output Prediction". *arXiv preprint arXiv:2502.07316*.
133. [△]Bajpai Y, Chopra B, Biyani P, Aslan C, Coleman D, Gulwani S, Parnin C, Radhakrishna A, Soares G. "Let's Fix this Together: Conversational Debugging with GitHub Copilot." In: *2024 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE; 2024. p. 1-12.
134. [△]Li R, Zhang Y, Yang D (2024). "Sketch2Code: Evaluating Vision-Language Models for Interactive Web Design Prototyping". *arXiv preprint arXiv:2410.16232*. Available from: <https://arxiv.org/abs/2410.16232>.
135. [△]DeepSeek-AI, Guo D, Yang D, Zhang H, Song J, Zhang R, Xu R, Zhu Q, Ma S, Wang P, Bi X, Zhang X, Yu X, Wu Y, Wu ZF, Gou Z, Shao Z, Li Z, Gao Z, Liu A, Xue B, Wang B, Wu B, Feng B, Lu C, Zhao C, Deng C, Zhang C, Ruan

C, Dai D, Chen D, Ji D, Li E, Lin F, Dai F, Luo F, Hao G, Chen G, Li G, Zhang H, Bao H, Xu H, Wang H, Ding H, Xin H, Gao H, Qu H, Li H, Guo J, Li J, Wang J, Chen J, Yuan J, Qiu J, Li J, Cai JL, Ni J, Liang J, Chen J, Dong K, Hu K, Gao K, Guan K, Huang K, Yu K, Wang L, Zhang L, Zhao L, Wang L, Zhang L, Xu L, Xia L, Zhang M, Zhang M, Tang M, Li M, Wang M, Li M, Tian N, Huang P, Zhang P, Wang Q, Chen Q, Du Q, Ge R, Zhang R, Pan R, Wang R, Chen RJ, Jin RL, Chen R, Lu S, Zhou S, Chen S, Ye S, Wang S, Yu S, Zhou S, Pan S, Li SS, Zhou S, Wu S, Ye S, Yun T, Pei T, Sun T, Wang T, Zeng W, Zhao W, Liu W, Liang W, Gao W, Yu W, Zhang W, Xiao WL, An W, Liu X, Wang X, Chen X, Nie X, Cheng X, Liu X, Xie X, Liu X, Yang X, Li X, Su X, Lin X, Li XQ, Jin X, Shen X, Chen X, Sun X, Wang X, Song X, Zhou X, Wang X, Shan X, Li YK, Wang YQ, Wei YX, Zhang Y, Xu Y, Li Y, Zhao Y, Sun Y, Wang Y, Yu Y, Zhang Y, Shi Y, Xiong Y, He Y, Piao Y, Wang Y, Tan Y, Ma Y, Liu Y, Guo Y, Ou Y, Wang Y, Gong Y, Zou Y, He Y, Xiong Y, Luo Y, You Y, Liu Y, Zhou Y, Zhu YX, Xu Y, Huang Y, Li Y, Zheng Y, Zhu Y, Ma Y, Tang Y, Zha Y, Yan Y, Ren Z, Ren Z, Sha Z, Fu Z, Xu Z, Xie Z, Zhang Z, Hao Z, Ma Z, Yan Z, Wu Z, Gu Z, Zhu Z, Liu Z, Li Z, Xie Z, Song Z, Pan Z, Huang Z, Xu Z, Zhang Z, Zhang Z. "DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning". arXiv. 2025. Available from: <https://arxiv.org/abs/2501.12948>.

136. ^aWei Y, Duchenne O, Copet J, Carbonneaux Q, Zhang L, Fried D, Synnaeve G, Singh R, Wang SI (2025). "SW E-RL: Advancing LLM Reasoning via Reinforcement Learning on Open Software Evolution". arXiv. [2502.18449](#).
137. ^aPan J, Wang X, Neubig G, Jaitly N, Ji H, Suhr A, Zhang Y (2024). "Training Software Engineering Agents and Verifiers with SWE-Gym". arXiv. [arXiv:2412.21139](#) [cs.SE].
138. ^ΔGuo X, Wang X, Chen Y, Li S, Han C, Li M, Ji H (2025). "SyncMind: Measuring Agent Out-of-Sync Recovery in Collaborative Software Engineering". arXiv preprint arXiv:2502.06994.
139. ^ΔXie Y, Xie A, Sheth D, Liu P, Fried D, Rose C (2025). "RepoST: Scalable Repository-Level Coding Environment Construction with Sandbox Testing". arXiv preprint arXiv:2503.07358. Available from: <https://arxiv.org/abs/2503.07358>.
140. ^ΔTeam K, Du A, Gao B, Xing B, Jiang C, Chen C, Li C, Xiao C, Du C, Liao C, et al. Kimi k1. 5: Scaling reinforcement learning with llms. arXiv preprint arXiv:2501.12599. 2025.
141. ^ΔSkalse J, Howe N, Krashennikov D, Krueger D (2022). "Defining and characterizing reward gaming". *Advances in Neural Information Processing Systems*. 35: 9460–9471.
142. ^ΔBaker B, Huizinga J, Gao L, Dou Z, Guan MY, Madry A, Zaremba W, Pachocki J, Farhi D (2025). "Monitoring reasoning models for misbehavior and the risks of promoting obfuscation". arXiv preprint arXiv:2503.11926.
143. ^ΔDenison C, MacDiarmid M, Barez F, Duvenaud D, Kravec S, Marks S, Schiefer N, Soklaski R, Tamkin A, Kaplan J, et al. Sycophancy to subterfuge: Investigating reward-tampering in large language models. arXiv preprint

rint arXiv:2406.10162. 2024.

144. [△]Papineni K, Roukos S, Ward T, Zhu WJ. "Bleu: a method for automatic evaluation of machine translation." In: *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*. 2002. p. 311–318.
145. [△]Ren S, Guo D, Lu S, Zhou L, Liu S, Tang D, Sundaresan N, Zhou M, Blanco A, Ma S (2020). "Codebleu: a method for automatic evaluation of code synthesis". arXiv preprint arXiv:2009.10297. [arXiv:2009.10297](https://arxiv.org/abs/2009.10297).
146. [△]Zhang T, Kishore V, Wu F, Weinberger KQ, Artzi Y (2019). "Bertscore: Evaluating text generation with bert". arXiv preprint arXiv:1904.09675.
147. [△]Zhou S, Alon U, Agarwal S, Neubig G (2023). "Codebertscore: Evaluating code generation with pretrained models of code". arXiv preprint arXiv:2302.05527.
148. [△]Maertens R, Van Petegem C, Strijbol N, Baeyens T, Jacobs AC, Dawyndt P, Mesuere B (2022). "Dolos: Language-agnostic plagiarism detection in source code". *Journal of Computer Assisted Learning*. 38 (4): 1046–1061.
149. [△]Ma Z, Peng C, Gao P, Meng X, Zou Y, Xie B (2025). "SoRFT: Issue Resolving with Subtask-oriented Reinforced Fine-Tuning". arXiv preprint arXiv:2502.20127. Available from: <https://arxiv.org/abs/2502.20127>.
150. [△]Akyürek E, Damani M, Qiu L, Guo H, Kim Y, Andreas J (2024). "The Surprising Effectiveness of Test-Time Training for Abstract Reasoning". arXiv. [arXiv:2411.07279](https://arxiv.org/abs/2411.07279).
151. [△]Sun Y, Wang X, Zhuang L, Miller J, Hardt M, Efros AA. Test-time training with self-supervision for generalization under distribution shifts. In: *ICML; 2020*.
152. [△]Lester B, Al-Rfou R, Constant N (2021). "The power of scale for parameter-efficient prompt tuning". arXiv preprint arXiv:2104.08691. [arXiv:2104.08691](https://arxiv.org/abs/2104.08691).
153. [△]Li XL, Liang P (2021). "Prefix-tuning: Optimizing continuous prompts for generation". arXiv preprint arXiv:2101.00190. Available from: <https://arxiv.org/abs/2101.00190>.
154. [△]Szegedy C. A promising path towards autoformalization and general artificial intelligence. In: *Intelligent Computer Mathematics: 13th International Conference, CICM 2020, Bertinoro, Italy, July 26–31, 2020, Proceedings 13*. Springer; 2020. p. 3–20.
155. [△]Endres M, Fakhoury S, Chakraborty S, Lahiri SK (2024). "Can large language models transform natural language intent into formal method postconditions?" *Proceedings of the ACM on Software Engineering*. 1 (FSE): 1889–1912.
156. [△]Lange RT, Prasad A, Sun Q, Faldor M, Tang Y, Ha D (2025). "The AI CUDA Engineer: Agentic CUDA Kernel Discovery, Optimization and Composition".

157. ^a Vijayvargiya S, Zhou X, Yerukola A, Sap M, Neubig G. Interactive agents to overcome ambiguity in software engineering. 2025. Available from: <https://arxiv.org/abs/2502.13069>.
158. ^ΔShao Y, Samuel V, Jiang Y, Yang J, Yang D (2024). "Collaborative Gym: A Framework for Enabling and Evaluating Human-Agent Collaboration". arXiv. Available from: <https://arxiv.org/abs/2412.15701>.
159. ^ΔZhao Y, Gong L, Zhang H, Yu Y, Huang Z (2023). "How to get better embeddings with code pre-trained models? An empirical study". arXiv preprint arXiv:2311.08066. Available from: <https://arxiv.org/abs/2311.08066>.
160. ^ΔNye M, Pu Y, Bowers M, Andreas J, Tenenbaum JB, Solar-Lezama A (2020). "Representing partial programs with blended abstract semantics". arXiv preprint arXiv:2012.12964.
161. ^ΔZohar A, Wolf L (2018). "Automatic program synthesis of long programs with a learned garbage collector". *Advances in neural information processing systems*. 31.
162. ^ΔEllis K, Nye M, Pu Y, Sosa F, Tenenbaum J, Solar-Lezama A (2019). "Write, execute, assess: Program synthesis with a repl". *Advances in Neural Information Processing Systems*. 32.
163. ^ΔChen X, Song D, Tian Y (2021). "Latent execution for neural program synthesis beyond domain-specific languages". *Advances in Neural Information Processing Systems*. 34: 22196–22208.
164. ^ΔIzacard G, Grave E (2020). "Leveraging passage retrieval with generative models for open domain question answering". arXiv preprint arXiv:2007.01282. Available from: <https://arxiv.org/abs/2007.01282>.
165. ^ΔBorgeaud S, Mensch A, Hoffmann J, Cai T, Rutherford E, Millican K, Van Den Driessche GB, Lespiau JB, Damoc B, Clark A, et al. Improving language models by retrieving from trillions of tokens. In: *International conference on machine learning*. PMLR; 2022. p. 2206–2240.
166. ^a Izacard G, Lewis P, Lomeli M, Hosseini L, Petroni F, Schick T, Dwivedi-Yu J, Joulin A, Riedel S, Grave E (2023). "Atlas: Few-shot learning with retrieval augmented language models". *Journal of Machine Learning Research*. 24 (251): 1–43.
167. ^ΔShi W, Min S, Yasunaga M, Seo M, James R, Lewis M, Zettlemoyer L, Yih W (2023). "Replug: Retrieval-augmented black-box language models". arXiv preprint arXiv:2301.12652.
168. ^ΔAsare O, Nagappan M, Asokan N (2023). "Is github's copilot as bad as humans at introducing vulnerabilities in code?" *Empirical Software Engineering*. 28 (6): 129.
169. ^ΔFu Y, Liang P, Tahir A, Li Z, Shahin M, Yu J, Chen J (2023). "Security weaknesses of copilot generated code in github". arXiv preprint arXiv:2310.02059.
170. ^ΔEl-Kishky A, Wei A, Saraiva A, Minaev B, Selsam D, Dohan D, Song F, Lightman H, Clavera I, Pachocki J, et al. Competitive programming with large reasoning models. arXiv preprint arXiv:2502.06807. 2025.

171. [△]Cousot P, Cousot R (1977). "Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints". In: *Proceedings of the 4th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*. pp. 238–252.
172. [△]Godefroid P, Klarlund N, Sen K (2005). "DART: Directed automated random testing". In: *Proceedings of the 2005 ACM SIGPLAN conference on Programming language design and implementation*. pp. 213–223.
173. [△]Sen K, Marinov D, Agha G (2005). "CUTE: A concolic unit testing engine for C". *ACM SIGSOFT Software Engineering Notes*. **30** (5): 263–272.
174. [△]Clarke EM. "Model checking." In: *Foundations of Software Technology and Theoretical Computer Science: 17th Conference Kharagpur, India, December 18–20, 1997 Proceedings* 17. Springer; 1997. p. 54--56.
175. [△]Cardelli L (1996). "Type systems". *ACM Computing Surveys (CSUR)*. **28** (1): 263–264.
176. [△]Gong L, Elhoushi M, Cheung A (2024). "AST-T5: Structure-Aware Pretraining for Code Generation and Understanding". *arXiv preprint arXiv:2401.03003*.
177. [△]Poesia G, Polozov A, Le V, Tiwari A, Soares G, Meek C, Gulwani S (2022). "Synchronesh: Reliable Code Generation from Pre-trained Language Models". *International Conference on Learning Representations*. Available from: <https://openreview.net/forum?id=KmtVD97J43e>.
178. [△]Geng S, Josifoski M, Peyrard M, West R (2023). "Grammar-constrained decoding for structured NLP tasks without finetuning". *arXiv preprint arXiv:2305.13971*. Available from: <https://arxiv.org/abs/2305.13971>.
179. [△]Wei Y, Xia CS, Zhang L (2023). "Copiloting the copilots: Fusing large language models with completion engines for automated program repair". In: *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. pp. 172–184.
180. [△]Burstall RM, Darlington J (1977). "A transformation system for developing recursive programs". *Journal of the ACM (JACM)*. **24** (1): 44–67.
181. [△]Puschel M, Moura JM, Johnson JR, Padua D, Veloso MM, Singer BW, Xiong J, Franchetti F, Gacic A, Voronenko Y, et al. SPIRAL: Code generation for DSP transforms. *Proceedings of the IEEE*. **93**(2):232–275, 2005.
182. [△]Ragan-Kelley J, Barnes C, Adams A, Paris S, Durand F, Amarasinghe S (2013). "Halide: a language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines". *Acm Sigplan Notices*. **48** (6): 519–530.
183. [△]Hong C, Bhatia S, Haan A, Dong SK, Nikiforov D, Cheung A, Shao YS. "LLM-Aided Compilation for Tensor Accelerators." In: *2024 IEEE LLM Aided Design Workshop (LAD)*. IEEE; 2024. p. 1-14.
184. [△]Al Madi N. How readable is model-generated code? examining readability and visual inspection of github copilot. In: *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*

- g, ASE '22, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450394758. doi:[10.1145/3551349.3560438](https://doi.org/10.1145/3551349.3560438).
185. [△]Singhal S, Kumar V. "Creating Thorough Tests for AI-Generated Code is Hard." In: *Proceedings of the 16th Annual ACM India Compute Conference, COMPUTE '23, Hyderabad, India*. New York, NY, USA: Association for Computing Machinery; 2023. p. 108–111. doi:[10.1145/3627217.3627238](https://doi.org/10.1145/3627217.3627238).
 186. [△]Jiang Y, Shao Y, Ma D, Semnani SJ, Lam MS (2024). "Into the unknown unknowns: Engaged human learning through participation in language model agent conversations". *arXiv preprint arXiv:2408.15232*.
 187. [△][‡]Sun W, Miao Y, Li Y, Zhang H, Fang C, Liu Y, Deng G, Liu Y, Chen Z (2024). "Source code summarization in the era of large language models". *arXiv preprint arXiv:2407.07959*. Available from: <https://arxiv.org/abs/2407.07959>.
 188. [△]Ferdowsi K, Huang R, James MB, Polikarpova N, Lerner S (2024). "Validating AI-generated code with live programming". In: *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 2024: 1–8.
 189. [△]Pu Y, Ellis K, Kryven M, Tenenbaum J, Solar-Lezama A (2020). "Program synthesis with pragmatic communication". *Advances in neural information processing systems*. 33: 13249–13259.
 190. [△]Radford A, Wu J, Child R, Luan D, Amodei D, Sutskever I (2019). "Language Models are Unsupervised Multitask Learners". *S2CID 160025533*.
 191. [△]Agrawal LA, Kanade A, Goyal N, Lahiri SK, Rajamani SK (2023). "Guiding Language Models of Code with Global Context using Monitors". *arXiv*. <https://arxiv.org/abs/2306.10763>.
 192. [△]Park K, Wang J, Berg-Kirkpatrick T, Polikarpova N, D'Antoni L (2024). "Grammar-Aligned Decoding". *arXiv*. <https://arxiv.org/abs/2405.21047>.
 193. [△]Sun J, Tian Y, Zhou W, Xu N, Hu Q, Gupta R, Wieting JF, Peng N, Ma X (2023). "Evaluating Large Language Models on Controlled Generation Tasks". *arXiv*. <https://arxiv.org/abs/2310.14542>.
 194. [△]Zettlemoyer LS, Collins M (2012). "Learning to Map Sentences to Logical Form: Structured Classification with Probabilistic Categorical Grammars". *arXiv*. Available from: <https://arxiv.org/abs/1207.1420>.
 195. [△]Wong YW, Mooney R. Learning for semantic parsing with statistical machine translation. In: Moore RC, Bilmes J, Chu-Carroll J, Sanderson M, editors. *Proceedings of the Human Language Technology Conference of the NAACL, Main Conference*. New York City, USA: Association for Computational Linguistics; 2006. p. 439–446.
 196. [△]Zhong V, Xiong C, Socher R (2017). "Seq2SQL: Generating Structured Queries from Natural Language using Reinforcement Learning". *arXiv. cs.CL*. Available from: [arXiv:1709.00103](https://arxiv.org/abs/1709.00103).

197. [△]Yu T, Zhang R, Yang K, Yasunaga M, Wang D, Li Z, Ma J, Li I, Yao Q, Roman S, Zhang Z, Radev D (2019). "Spider: A Large-Scale Human-Labeled Dataset for Complex and Cross-Domain Semantic Parsing and Text-to-SQL Task". *arXiv. cs.CL*. Available from: <https://arxiv.org/abs/1809.08887>.
198. [△]OpenAI R (2023). "GPT-4 technical report. *arXiv 2303.08774*". View in Article.
199. [△]Du X, Liu M, Wang K, Wang H, Liu J, Chen Y, Feng J, Sha C, Peng X, Lou Y (2023). "ClassEval: A Manually-Crafted Benchmark for Evaluating LLMs on Class-level Code Generation". *arXiv. cs.CL*. Available from: <https://arxiv.org/abs/2308.01861>.
200. [△]Cao J, Chen Z, Wu J, Cheung SC, Xu C (2024). "JavaBench: A Benchmark of Object-Oriented Code Generation for Evaluating Large Language Models". *arXiv. 2406.12902*.
201. [△]Wang S, Ding L, Shen L, Luo Y, Du B, Tao D (2024). "OOP: Object-Oriented Programming Evaluation Benchmark for Large Language Models". *arXiv*. Available from: <https://arxiv.org/abs/2401.06628>.
202. [△]OpenAI (2023). "GPT-4 Demo: From Sketch to Website". YouTube. Available from: <https://www.youtube.com/watch?v=outcGtbnMuQ>. Accessed: 2025-03-26.
203. [△]Liu V, Kazi RH, Wei L-Y, Fisher M, Langlois T, Walker S, Chilton L (2025). "LogoMotion: Visually-Grounded Code Synthesis for Creating and Editing Animation". *arXiv*. Available from: <https://arxiv.org/abs/2405.07065>.
204. [△]Nandi C, Willsey M, Anderson A, Wilcox JR, Darulova E, Grossman D, Tatlock Z (2020). "Synthesizing structured CAD models with equality saturation and inverse transformations". In: *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2020*. New York, NY, USA: Association for Computing Machinery. p. 31–44. ISBN 9781450376136.
205. [△]Qiu Z, Liu W, Feng H, Liu Z, Xiao TZ, Collins KM, Tenenbaum JB, Weller A, Black MJ, Schölkopf B (2024). "Can large language models understand symbolic graphics programs?" *arXiv. https://arxiv.org/abs/2408.08313*.
206. [△][△]Cassano F, Gouwar J, Nguyen D, Nguyen S, Phipps-Costin L, Pinckney D, Yee M-H, Zi Y, Anderson CJ, Feldman MQ, et al. MultiPL-E: a scalable and polyglot approach to benchmarking neural code generation. *IEEE Transactions on Software Engineering. 49(7):3675–3691, 2023*.
207. [△]Chai L, Liu S, Yang J, Yin Y, Jin K, Liu J, Sun T, Zhang G, Ren C, Guo H, et al. Mceval: Massively multilingual code evaluation. *arXiv preprint arXiv:2406.07436. 2024*.
208. [△]Liu M, Pinckney N, Khailany B, Ren H (2023). "VerilogEval: Evaluating large language models for verilog code generation". In: *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*. IEEE. pp. 1–8.

209. [△]Florath A (2024). "Enhancing formal theorem proving: a comprehensive dataset for training AI models on Coq code". arXiv preprint arXiv:2403.12627. Available from: <https://arxiv.org/abs/2403.12627>.
210. [△]Pei Z, Zhen HL, Yuan M, Huang Y, Yu B (2024). "Betternv: Controlled verilog generation with discriminative guidance". arXiv preprint arXiv:2402.03375.
211. [△]Paul I, Glava_v G, Gurevych I (2024). "Ircoder: Intermediate representations make language models robust multilingual code generators". arXiv preprint arXiv:2403.03894. Available from: <https://arxiv.org/abs/2403.03894>.
212. [△]Cassano F, Gouwar J, Lucchetti F, Schlesinger C, Freeman A, Anderson CJ, Feldman MQ, Greenberg M, Jangda A, Guha A (2024). "Knowledge transfer from high-resource to low-resource programming languages for code llms". *Proceedings of the ACM on Programming Languages*. 8 (OOPSLA2): 677–708.
213. [△]Orlanski G, Xiao K, Garcia X, Hui J, Howland J, Malmaud J, Austin J, Singh R, Catasta M. Measuring the impact of programming language distribution. In: *International Conference on Machine Learning*. PMLR; 2023. p. 26619–26645.
214. [△]Yang K, Swope A, Gu A, Chalamala R, Song P, Yu S, Godil S, Prenger RJ, Anandkumar A (2023). "Leandojo: Theorem proving with retrieval-augmented language models". *Advances in Neural Information Processing Systems*. 36: 21573–21612.
215. [△]Zhou S, Alon U, Xu FF, Wang Z, Jiang Z, Neubig G (2022). "Docprompting: Generating code by retrieving the docs". arXiv preprint arXiv: 2207.05987.
216. [△]Vero M, Mündler N, Chibotaru V, Raychev V, Baader M, Jovanović N, He J, Vechev M (2025). "BaxBench: Can LLMs Generate Correct and Secure Backends?" arXiv. <https://arxiv.org/abs/2502.11844>.
217. [△]Siddiq ML, Santos JCS. SecurityEval dataset: mining vulnerability examples to evaluate machine learning-based code generation techniques. In: *Association for Computing Machinery*; 2022. p. 29–33. ISBN 9781450394574.
218. [△]He J, Vero M, Krasnopolka G, Vechev M (2024). "Instruction Tuning for Secure Code Generation". arXiv. <https://arxiv.org/abs/2402.09497>.
219. [△]Hajipour H, Hassler K, Holz T, Schfönherr L, Fritz M (2023). "CodeLMSec Benchmark: Systematically Evaluating and Finding Security Vulnerabilities in Black-Box Code Language Models". arXiv. Available from: <https://arxiv.org/abs/2302.04012>.
220. [△]Peng J, Cui L, Huang K, Yang J, Ray B (2025). "CWEval: Outcome-driven Evaluation on Functionality and Security of LLM Code Generation". arXiv. Available from: <https://arxiv.org/abs/2501.08200>.

221. [△]Bhatt M, Chennabasappa S, Nikolaidis C, Wan S, Evtimov I, Gabi D, Song D, Ahmad F, Aschermann C, Fontana L, Frolov S, Giri RP, Kapil D, Kozyrakis Y, LeBlanc D, Milazzo J, Straumann A, Synnaeve G, Vontimitta V, Whitman S, Saxe J (2023). "Purple Llama CyberSecEval: A Secure Coding Benchmark for Language Models". *arXiv*. Available from: <https://arxiv.org/abs/2312.04724>.
222. [△]Wan S, Nikolaidis C, Song D, Molnar D, Crnkovich J, Grace J, Bhatt M, Chennabasappa S, Whitman S, Ding S, et al. Cyberseceval 3: Advancing the evaluation of cybersecurity risks and capabilities in large language models. *arXiv preprint arXiv:2408.01605*. 2024.
223. [△]Ma J, Sreedhar K, Liu V, Perez PA, Wang S, Sahni R, Chilton LB (2025). "DynEx: Dynamic Code Synthesis with Structured Design Exploration for Accelerated Exploratory Programming". *arXiv*. Available from: <https://arxiv.org/abs/2410.00400>.
224. [△]Champa AI, Rabbi MF, Nachuma C, Zibran MF. "ChatGPT in Action: Analyzing Its Use in Software Development." In: 2024 IEEE/ACM 21st International Conference on Mining Software Repositories (MSR), 2024. p. 182-186.
225. [△]Xiao T, Treude C, Hata H, Matsumoto K (2024). "DevGPT: Studying developer-chatgpt conversations". In: *Proceedings of the 21st International Conference on Mining Software Repositories*. pp. 227–230.
226. [△]Surameery NM, Shakor MY (2023). "Use chat gpt to solve programming bugs". *International Journal of Information Technology and Computer Engineering*. **31**: 17–22.
227. [△]Kazemitabaar M, Chow J, Ma CK, Ericson BJ, Weintrop D, Grossman T (2023). "Studying the effect of AI code generators on supporting novice learners in introductory programming". In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 2023: 1–23.
228. [△]Kazemitabaar M, Hou X, Henley A, Ericson BJ, Weintrop D, Grossman T (2023). "How novices use LLM-based code generators to solve CS1 coding tasks in a self-paced learning environment". *Proceedings of the 23rd Koli calling international conference on computing education research*. pages 1–12.
229. [△]Prather J, Reeves BN, Denny P, Becker BA, Leinonen J, Luxton-Reilly A, Powell G, Finnie-Ansley J, Santos EA (2023). "It's weird that it knows what I want": Usability and interactions with copilot for novice programmers. *ACM Transactions on Computer-Human Interaction*. **31** (1): 1–31.
230. [△]Sheese B, Liffiton M, Savelka J, Denny P (2024). "Patterns of student help-seeking when using a large language model-powered programming assistant." In: *Proceedings of the 26th Australasian computing education conference*. pp. 49–57.
231. [△]Yan H, Latoza TD, Yao Z (2024). "IntelliExplain: Enhancing Conversational Code Generation for Non-Professional Programmers". *arXiv preprint arXiv:2405.10250*.

232. [△]Lahiri SK, Fakhoury S, Naik A, Sakkas G, Chakraborty S, Musuvathi M, Choudhury P, von Veh C, Inala JP, Wang C, et al. Interactive code generation via test-driven user-intent formalization. *arXiv preprint arXiv:2208.05950*. 2022.
233. [△]Fakhoury S, Naik A, Sakkas G, Chakraborty S, Lahiri SK (2024). "Llm-based test-driven interactive code generation: User study and empirical evaluation". *IEEE Transactions on Software Engineering*. IEEE.
234. [△]Pailoor S, Wang Y, Dillig I (2024). "Semantic Code Refactoring for Abstract Data Types". *Proc. ACM Program. Lang.*. **POPL**.
235. [△]Islam M, Jha AK, Nadi S, Akhmetov I (2023). "PyMigBench: A Benchmark for Python Library Migration". In: *2023 IEEE/ACM 20th International Conference on Mining Software Repositories (MSR)*.
236. [△]Omidvar Tehrani B, Anubhai A (2024). "Evaluating Human-AI Partnership for LLM-based Code Migration." In: *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, pp. 1–8.
237. [△]Eniser HF, Zhang H, David C, Wang M, Christakis M, Paulsen B, Dodds J, Kroening D (2024). "Towards Translating Real-World Code with LLMs: A Study of Translating to Rust". *arXiv*. [arXiv:2405.11514 \[cs.SE\]](https://arxiv.org/abs/2405.11514).
238. [△]Li R, Wang B, Li T, Saxena P, Kundu A (2025). "Translating C to Rust: Lessons from a User Study". In: *Proceedings 2025 Network and Distributed System Security Symposium, NDSS 2025*. Internet Society, 2025.
239. [△]Mankowitz DJ, Michi A, Zhernov A, Gelmi M, Selvi M, Paduraru C, Leurent E, Iqbal S, Lespiau JB, Ahern A, et al. Faster sorting algorithms discovered using deep reinforcement learning. *Nature*. **618**(7964):257–263, 2023.
240. [△]Gong J, Voskanyan V, Brookes P, Wu F, Jie W, Xu J, Giavrimis R, Basios M, Kanthan L, Wang Z (2025). "Language models for code optimization: Survey, challenges and future directions". *arXiv preprint arXiv:2501.01277*.
241. [△]Lemieux C, Inala JP, Lahiri SK, Sen S (2023). "Codamosa: Escaping coverage plateaus in test generation with pre-trained large language models". In: *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE. pp. 919–931.
242. [△]Vikram V, Lemieux C, Sunshine J, Padhye R (2023). "Can large language models write good property-based tests?" *arXiv preprint arXiv:2307.04346*.
243. [△]Li K, Yuan Y (2024). "Large Language Models as Test Case Generators: Performance Evaluation and Enhancement". *arXiv*. [arXiv:2404.13340](https://arxiv.org/abs/2404.13340).
244. [△]Mündler N, Müller M, He J, Vechev M (2025). "SWT-bench: Testing and validating real-world bug-fixes with code agents". *Advances in Neural Information Processing Systems*. **37**: 81857–81887.

245. [△]Ryan G, Jain S, Shang M, Wang S, Ma X, Ramanathan MK, Ray B (2024). "Code-aware prompting: A study of coverage-guided test generation in regression setting using llm". *Proceedings of the ACM on Software Engineering*. **1** (FSE): 951–971.
246. [△]Chen B, Zhang F, Nguyen A, Zan D, Lin Z, Lou JG, Chen W (2022). "Codet: Code generation with generated tests". *arXiv preprint arXiv:2207.10397*. [arXiv:2207.10397](https://arxiv.org/abs/2207.10397).
247. [△]Zhang K, Wang D, Xia J, Wang WY, Li L (2023). "ALGO: Synthesizing Algorithmic Programs with Generated Oracle Verifiers". *arXiv preprint arXiv:2305.14591*. [arXiv:2305.14591](https://arxiv.org/abs/2305.14591).
248. [△]Chen X, Tao Z, Zhang K, Zhou C, Gu W, He Y, Zhang M, Cai X, Zhao H, Jin Z (2025). "Revisit self-debugging with self-generated tests for code generation". *arXiv preprint arXiv:2501.12793*.
249. [△]Miller BP, Fredriksen L, So B (1990). "An empirical study of the reliability of UNIX utilities". *Communications of the ACM*. **33** (12): 32–44.
250. [△]Xia CS, Paltenghi M, Le Tian J, Pradel M, Zhang L (2024). "Fuzz4all: Universal fuzzing with large language models". In: *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. pp. 1–13.
251. [△]Yang C, Zhao Z, Zhang L (2023). "Kernelgpt: Enhanced kernel fuzzing via large language models". *arXiv preprint arXiv:2401.00563*. Available from: <https://arxiv.org/abs/2401.00563>.
252. [△]Liu D, Metzman J, Chang O, Google Open Source Security Team (2023). "AI-Powered Fuzzing: Breaking the Bug Hunting Barrier". <https://security.googleblog.com/2023/08/ai-powered-fuzzing-breaking-bug-hunting.html>.
253. [△]Liu D, Chang O, Metzman J, Sablotny M, Maruseac M. OSS-Fuzz-Gen: Automated Fuzz Target Generation [software]. May 2024. Available from: <https://github.com/google/oss-fuzz-gen>. Version: v1.0. License: Apache-2.0.
254. [△]Zhou Y, Liu S, Siow J, Du X, Liu Y (2019). "Devign: Effective Vulnerability Identification by Learning Comprehensive Program Semantics via Graph Neural Networks". In: *Neural Information Processing Systems*.
255. [△]Chakraborty S, Krishna R, Ding Y, Ray B (2020). "Deep Learning Based Vulnerability Detection: Are We There Yet?" *IEEE Transactions on Software Engineering*. **48**: 3280–3296.
256. [△]Dinella E, Dai H, Li Z, Naik M, Song L, Wang K. Hoppity: Learning graph transformations to detect and fix bugs in programs. In: *International Conference on Learning Representations*; 2020.
257. [△]Hin D, Kan A, Chen H, Babar MA. "LineVD: Statement-level Vulnerability Detection using Graph Neural Networks". In: *International Conference on Mining Software Repositories*, 2022.
258. [△]Li Y, Wang S, Nguyen TN. "Vulnerability detection with fine-grained interpretations." In: *Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*,

2021.

259. [△]Fu M, Tantithamthavorn C. "LineVul: A Transformer-based Line-Level Vulnerability Prediction". In: *International Conference on Mining Software Repositories*, 2022.
260. [△]Steenhoek B, Gao H, Le W (2023). "Dataflow Analysis-Inspired Deep Learning for Efficient Vulnerability Detection". arXiv preprint arXiv:2212.08108. [arXiv:2212.08108](https://arxiv.org/abs/2212.08108).
261. [△]Cheng X, Zhang G, Wang H, Sui Y (2022). "Path-sensitive code embedding via contrastive learning for software vulnerability detection". In: *International Symposium on Software Testing and Analysis*.
262. [△]Steenhoek B, Rahman MM, Roy MK, Alam MS, Barr ET, Le W (2024). "A Comprehensive Study of the Capabilities of Large Language Models for Vulnerability Detection". arXiv preprint arXiv:2403.17218.
263. [△]Ding Y, Fu Y, Ibrahim O, Sitawarin C, Chen X, Alomair B, Wagner D, Ray B, Chen Y (2024). "Vulnerability detection with code language models: How far are we?" arXiv preprint arXiv:2403.18624.
264. [△]Khare A, Dutta S, Li Z, Solko-Breslin A, Alur R, Naik M (2023). "Understanding the Effectiveness of Large Language Models in Detecting Security Vulnerabilities". arXiv preprint arXiv:2311.16169. 2023. Available from: <https://arxiv.org/abs/2311.16169>.
265. [△]Li H, Hao Y, Zhai Y, Qian Z (2024). "Enhancing static analysis for practical bug detection: An LLM-integrated approach". *Proc. ACM Program. Lang.* 8 (OOPSLA1).
266. [△]Li Z, Dutta S, Naik M (2024). "LLM-Assisted Static Analysis for Detecting Security Vulnerabilities". arXiv preprint arXiv:2405.17238. Available from: <https://arxiv.org/abs/2405.17238>.
267. [△]Wang C, Zhang W, Su Z, Xu X, Xie X, Zhang X (2024). "LLMDFA: Analyzing Dataflow in Code with Large Language Models". arXiv. [arXiv:2402.10754](https://arxiv.org/abs/2402.10754).
268. [△]Wang C, Liu J, Peng X, Liu Y, Lou Y (2023). "Boosting Static Resource Leak Detection via LLM-based Resource-Oriented Intention Inference". arXiv preprint arXiv:2311.04448.
269. [△]Ernst MD, Perkins JH, Guo PJ, McCamant S, Pacheco C, Tschantz MS, Xiao C (2007). "The Daikon system for dynamic detection of likely invariants". *Science of Computer Programming*. 69 (1-3): 35–45.
270. [△]Padon O, Immerman N, Shoham S, Karbyshev A, Sagiv M (2016). "Decidability of inferring inductive invariants". *ACM SIGPLAN Notices*. 51 (1): 217–231.
271. [△]Hangal S, Lam MS (2002). "Tracking down software bugs using automatic anomaly detection". In: *Proceedings of the 24th international conference on Software engineering*. pp. 291–301.
272. [△]Ernst MD, Cockrell J, Griswold WG, Notkin D (1999). "Dynamically discovering likely program invariants to support program evolution". In: *Proceedings of the 21st international conference on Software engineering*. p. 213–224.

273. [△]Zheng D, Sen K (2024). "Dynamic Inference of Likely Symbolic Tensor Shapes in Python Machine Learning Programs". In: *Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Practice*. pp. 147–156.
274. [△]Roychoudhury A, Pasareanu C, Pradel M, Ray B (2025). "AI Software Engineer: Programming with Trust". arXiv. Available from: <https://arxiv.org/abs/2502.13767>.
275. [△]Dinella E, Lahiri SK, Naik M (2024). "Inferring natural preconditions via program transformation." In: *Companion Proceedings of the 32nd ACM International Conference on the Foundations of Software Engineering, FSE 2024, New York, NY, USA: Association for Computing Machinery*. p. 657–658. ISBN 9798400706585.
276. [△]Ruan H, Zhang Y, Roychoudhury A (2024). "SpecRover: Code Intent Extraction via LLMs". arXiv. <https://arxiv.org/abs/2408.02232>.
277. [△]Dinella E, Lahiri S, Naik M (2024). "Program Structure Aware Precondition Generation". arXiv. Available from: <https://arxiv.org/abs/2310.02154>.
278. [△]Si X, Dai H, Raghothaman M, Naik M, Song L (2018). "Learning loop invariants for program verification". *Advances in Neural Information Processing Systems*. 31.
279. [△]Kamath A, Senthilnathan A, Chakraborty S, Deligiannis P, Lahiri SK, Lal A, Rastogi A, Roy S, Sharma R (2023). "Finding inductive loop invariants using large language models". arXiv preprint arXiv:2311.07948. [arXiv:2311.07948](https://arxiv.org/abs/2311.07948).
280. [△]Chakraborty S, Lahiri SK, Fakhoury S, Musuvathi M, Lal A, Rastogi A, Senthilnathan A, Sharma R, Swamy N (2023). "Ranking llm-generated loop invariants for program verification". arXiv preprint arXiv:2310.09342. [arXiv:2310.09342](https://arxiv.org/abs/2310.09342).
281. [△]Yu S, Wang T, Wang J (2023). "Loop Invariant Inference through SMT Solving Enhanced Reinforcement Learning". In: *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis*. pp. 175–187.
282. [△]Wei J, Durrett G, Dillig I (2023). "Typet5: Seq2seq type inference using static analysis". arXiv preprint arXiv:2303.09564. [arXiv:2303.09564](https://arxiv.org/abs/2303.09564).
283. [△]Peng Y, Wang C, Wang W, Gao C, Lyu MR. "Generative type inference for python." In: *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE; 2023. p. 988–999.
284. [△]Wang C, Zhang J, Lou Y, Liu M, Sun W, Liu Y, Peng X (2024). "Tiger: A generating-then-ranking framework for practical python type inference". arXiv preprint arXiv:2407.02095.
285. [△]Liu C, Wu X, Feng Y, Cao Q, Yan J (2024). "Towards General Loop Invariant Generation: A Benchmark of Programs with Memory Manipulation". *The Thirty-eight Conference on Neural Information Processing Systems*.

286. [△]Pei K, Guan J, Broughton M, Chen Z, Yao S, Williams-King D, Ummadisetty V, Yang J, Ray B, Jana S (2021). "StateFormer: fine-grained type recovery from binaries using generative state modeling". In: *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2021*. Association for Computing Machinery. ISBN 9781450385626.
287. [△]Zhu C, Li Z, Xue A, Bajaj AP, Gibbs W, Liu Y, Alur R, Bao T, Dai H, Doupé A, Naik M, Shoshitaishvili Y, Wang R, Machiry A. "TYGR: Type inference on stripped binaries using graph neural networks." In: *33rd USENIX Security Symposium (USENIX Security 24)*; 2024. ISBN 978-1-939133-44-1.
288. [△]Liu P, Sun J, Chen L, Yan Z, Zhang P, Sun D, Wang D, Li D (2025). "Control Flow-Augmented Decompiler based on Large Language Model". *arXiv*. [arXiv:2503.07215](https://arxiv.org/abs/2503.07215) [cs.SE].
289. [△]Liu P, Sun C, Zheng Y, Feng X, Qin C, Wang Y, Li Z, Sun L (2023). "Harnessing the power of llm to support binary taint analysis". *arXiv preprint arXiv:2310.08275*. [arXiv:2310.08275](https://arxiv.org/abs/2310.08275).
290. [△]Jin X, Larson J, Yang W, Lin Z (2023). "Binary Code Summarization: Benchmarking ChatGPT/GPT-4 and Other Large Language Models". *arXiv*. [arXiv:2312.09601](https://arxiv.org/abs/2312.09601).
291. [△]Liu Y, Meng R, Joty S, Savarese S, Xiong C, Zhou Y, Yavuz S (2024). "CodeXEmbed: A Generalist Embedding Model Family for Multilingual and Multi-task Code Retrieval". *arXiv*. [arXiv:2411.12644](https://arxiv.org/abs/2411.12644) [cs.SE].
292. [△]Avgustinov P, de Moor O, Jones Peyton M, Schafer M (2016). "QL: object-oriented queries on relational data". *ECOOP*.
293. [△][‡]Bouzenia I, Devanbu P, Pradel M (2024). "Repairagent: An autonomous, llm-based agent for program repair". *arXiv preprint arXiv:2403.17134*.
294. [△][‡]Xia CS, Deng Y, Dunn S, Zhang L (2024). "Agentless: Demystifying llm-based software engineering agents". *arXiv preprint arXiv:2407.01489*.
295. [△]Su CY, McMillan C (2024). "Distilled GPT for source code summarization". *Automated Software Engineering*. **31** (1): 22.
296. [△]Haldar R, Hockenmaier J (2024). "Analyzing the performance of large language models on code summarization". *arXiv preprint arXiv:2404.08018*.
297. [△]Ahmed T, Pai KS, Devanbu P, Barr E. "Automatic semantic augmentation of language model prompts (for code summarization)". In: *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. 2024. p. 1–13.
298. [△]Luo Q, Ye Y, Liang S, Zhang Z, Qin Y, Lu Y, Wu Y, Cong X, Lin Y, Zhang Y, et al. Repoagent: An llm-powered open-source framework for repository-level code documentation generation. *arXiv preprint arXiv:2402.1666*

7. 2024.

299. [^]Shi K, Alt[^]31nbfcken D, Anand S, Christodorescu M, Grfcfcnwedel K, Koenings A, Naidu S, Pathak A, Rasi M, Ribeiro F, et al. Natural language outlines for code: Literate programming in the llm era. arXiv preprint arXiv:2408.04820. 2024.
300. [^]Diggs C, Doyle M, Madan A, Scott S, Escamilla E, Zimmer J, Nekoo N, Ursino P, Bartholf M, Robin Z, et al. Leveraging LLMs for Legacy Code Modernization: Challenges and Opportunities for LLM-Generated Documentation. arXiv preprint arXiv:2411.14971. 2024.
301. [^]Sun T, Xu J, Li Y, Yan Z, Zhang G, Xie L, Geng L, Wang Z, Chen Y, Lin Q, et al. BitsAI-CR: Automated Code Review via LLM in Practice. arXiv preprint arXiv:2501.15134. 2025.
302. [^]Tufano R, Pascarella L, Tufano M, Poshyvanyk D, Bavota G (2021). "Towards automating code review activities". In: 2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE). IEEE. p. 163–174.
303. [^]Tufano R, Masiero S, Mastropaolo A, Pascarella L, Poshyvanyk D, Bavota G (2022). "Using pre-trained models to boost code review automation". Proceedings of the 44th international conference on software engineering. 2291–2302.
304. [^]Li Z, Lu S, Guo D, Duan N, Jannu S, Jenks G, Majumder D, Green J, Svyatkovskiy A, Fu S, et al. Automating code review activities by large-scale pre-training. In: Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering. 2022. p. 1035–1047.
305. [^]Li K, Zhu A, Zhao P, Song J, Liu J (2024). "Utilizing deep learning to optimize software development processes". arXiv preprint arXiv:2404.13630. Available from: <https://arxiv.org/abs/2404.13630>.
306. [^]Zhang Y, Ruan H, Fan Z, Roychoudhury A. "Autocoderover: Autonomous program improvement." In: Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis. 2024. p. 1592–1604.
307. [^]Tan SH, Dong Z, Gao X, Roychoudhury A (2018). "Repairing crashes in android apps". In: Proceedings of the 40th International Conference on Software Engineering, pp. 187–198.
308. [^]Just R, Jalali D, Ernst MD (2014). "Defects4J: A database of existing faults to enable controlled testing studies for Java programs". In: Proceedings of the 2014 international symposium on software testing and analysis, pp. 437–440.
309. [^]Silva A, Saavedra N, Monperrus M. "Gitbug-java: A reproducible benchmark of recent java bugs". In: Proceedings of the 21st International Conference on Mining Software Repositories. 2024. p. 118–122.
310. [^]Jiang Y, Liu H, Niu N, Zhang L, Hu Y (2021). "Extracting Concise Bug-Fixing Patches from Human-Written Patches in Version Control Systems". In: IEEE/ACM 43rd International Conference on Software Engineering (I

CSE 2021), Los Alamitos, CA, USA: IEEE Computer Society; 2021. p. 686–698. doi:[10.1109/ICSE43902.2021.00069](https://doi.org/10.1109/ICSE43902.2021.00069).

311. [△]Jiang Y, Liu H, Luo X, Zhu Z, Chi X, Niu N, Zhang Y, Hu Y, Bian P, Zhang L (2022). "BugBuilder: An Automated Approach to Building Bug Repository". *IEEE Transactions on Software Engineering*. pages 1–22. doi:[10.1109/TSE.2022.3177713](https://doi.org/10.1109/TSE.2022.3177713).
312. [△]Jiang Y, Liu H, Zhang Y, Ji W, Zhong H, Zhang L. Do bugs lead to unnaturalness of source code? In: *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2022*; 2022. p. 1085–1096. doi:[10.1145/3540250.3549149](https://doi.org/10.1145/3540250.3549149).
313. [△]Widyasari R, Sim SQ, Lok C, Qi H, Phan J, Tay Q, Tan C, Wee F, Tan JE, Yieh Y, et al. Bugsinpy: a database of existing bugs in python programs to enable controlled testing and debugging studies. In: *Proceedings of the 28th ACM joint meeting on european software engineering conference and symposium on the foundations of software engineering*. 2020. p. 1556–1560.
314. [△]Tomassi DA, Dmeiri N, Wang Y, Bhowmick A, Liu YC, Devanbu PT, Vasilescu B, Rubio-González C. "Bugswarm: Mining and continuously growing a dataset of reproducible failures and fixes." In: *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE; 2019. p. 339–349.
315. [△]Hu X, Kuang K, Sun J, Yang H, Wu F (2024). "Leveraging print debugging to improve code generation in large language models". *arXiv preprint arXiv:2401.05319*.
316. [△]Tan SH, Yi J, Mehtaev S, Roychoudhury A, et al. Codeflaws: a programming competition benchmark for evaluating automated program repair tools. In: *2017 IEEE/ACM 39th International Conference on Software Engineering Companion (ICSE-C)*. IEEE; 2017. p. 180–182.
317. [△]Zhang Q, Fang C, Xie Y, Ma Y, Sun W, Yang Y, Chen Z (2024). "A systematic literature review on large language models for automated program repair". *arXiv preprint arXiv:2405.01466*.
318. [△]Lee C, Xia CS, Yang L, Huang J, Zhu Z, Zhang L, Lyu MR (2024). "A unified debugging approach via llm-based multi-agent synergy". *arXiv preprint arXiv:2404.17153*.
319. [△]Madaan A, Tandon N, Gupta P, Hallinan S, Gao L, Wiegrefe S, Alon U, Dziri N, Prabhume S, Yang Y, et al. Self-refine: Iterative refinement with self-feedback. *arXiv preprint arXiv:2303.17651*. 2023.
320. [△]Chen X, Lin M, Schärli N, Zhou D. "Teaching large language models to self-debug". In: *International Conference on Learning Representations (ICLR)*; 2024.
321. [△]Zhang K, Li Z, Li J, Li G, Jin Z. "Self-Edit: Fault-Aware Code Editor for Code Generation." In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Toronto, Canada: Association for Computational Linguistics; 2023. p. 769–787.

322. [△]Olausson TX, Inala JP, Wang C, Gao J, Solar-Lezama A. "Is Self-Repair a Silver Bullet for Code Generation?" In: *International Conference on Learning Representations (ICLR)*; 2024.
323. [△]Zhong L, Wang Z, Shang J (2024). "Debug like a human: A large language model debugger via verifying runtime execution step-by-step". *arXiv preprint arXiv:2402.16906*.
324. [△]Tang H, Hu K, Zhou J, Zhong SC, Zheng WL, Si X, Ellis K (2025). "Code repair with llms gives an exploration-exploitation tradeoff". *Advances in Neural Information Processing Systems*. 37: 117954–117996.
325. [△]Lu S, Guo D, Ren S, Huang J, Svyatkovskiy A, Blanco A, Clement C, Drain D, Jiang D, Tang D, et al. (2021). "Codexglue: A machine learning benchmark dataset for code understanding and generation". *arXiv preprint arXiv:2102.04664*. [arXiv:2102.04664](https://arxiv.org/abs/2102.04664).
326. [△]Nam D, Macvean A, Hellendoorn V, Vasilescu B, Myers B (2024). "Using an llm to help with code understanding". In: *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. pp. 1–13.
327. [△]Yang D, Liu T, Zhang D, Simoulin A, Liu X, Cao Y, Teng Z, Qian X, Yang G, Luo J, et al. Code to think, think to code: A survey on code-enhanced reasoning and reasoning-driven code intelligence in LLMs. *arXiv preprint arXiv:2502.19411*. 2025.
328. [△]Chaudhary D, Vadlamani SL, Thomas D, Nejati S, Sabetzadeh M (2024). "Developing a Llama-Based Chatbot for CI/CD Question Answering: A Case Study at Ericsson". *arXiv*. Available from: <https://arxiv.org/abs/2408.09277>.
329. [△]Bouzenia I, Pradel M (2024). "You Name It, I Run It: An LLM Agent to Execute Tests of Arbitrary Projects". *arXiv*. [arXiv:2412.10133](https://arxiv.org/abs/2412.10133).
330. [△]Yin Z, Ma X, Zheng J, Zhou Y, Bairavasundaram LN, Pasupathy S (2011). "An empirical study on configuration errors in commercial and open source systems". In: *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles, SOSP '11*, New York, NY, USA: Association for Computing Machinery. p. 159–172. ISBN 9781450309776.
331. [△]Verdet A, Hamdaqa M, Da Silva L, Khomh F (2023). "Exploring Security Practices in Infrastructure as Code: An Empirical Study". *arXiv*. [arXiv:2308.03952](https://arxiv.org/abs/2308.03952).
332. [△]Lamport L (1994). "Introduction to TLA".
333. [△]Leino KRM. "Dafny: An automatic program verifier for functional correctness." In: *International conference on logic for programming artificial intelligence and reasoning*. Springer; 2010. p. 348–370.
334. [△]Nipkow T, Wenzel M, Paulson LC. Isabelle/HOL: a proof assistant for higher-order logic. Springer; 2002.
335. [△]Lattuada A, Hance T, Bosamiya J, Brun M, Cho C, LeBlanc H, Srinivasan P, Achermann R, Chajed T, Hawblitzel C, et al. Verus: A practical foundation for systems verification. In: *Proceedings of the ACM SIGOPS 30th S*

ymposium on Operating Systems Principles. 2024. p. 438–454.

336. ^ΔCousot P, Cousot R, Feret J, Mauborgne L, Miné A, Monniaux D, Rival X. The ASTRÉE analyzer. In: *Programming Languages and Systems: 14th European Symposium on Programming, ESOP 2005, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2005, Edinburgh, UK, April 4–8, 2005. Proceedings 14*. Springer; 2005. p. 21–30.
337. ^ΔErbsen A, Philipoom J, Jamner D, Lin A, Gruetter S, Pit-Claudel C, Chlipala A (2024). "Foundational integration verification of a cryptographic server". *Proceedings of the ACM on Programming Languages*. 8 (PLDI): 1704–1729.
338. ^ΔErbsen A, Gruetter S, Choi J, Wood C, Chlipala A. Integration verification across software and hardware for a simple embedded system. In: *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. 2021. p. 604–619.
339. ^ΔKlein G, Elphinstone K, Heiser G, Andronick J, Cock D, Derrin P, Elkaduwe D, Engelhardt K, Kolanski R, Norrish M, et al. (2009). "seL4: Formal verification of an OS kernel." In *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*, pages 207–220.
340. ^ΔGoel S, Keizer A, Siddharth, Peng Y, Morrison K, Wetzler N, de Moura L, Ebeid N, Lee J, Letson A, Cicolini L, Kong S (2024). *leanprover/LNSym*. Available from: <https://github.com/leanprover/LNSym>.
341. ^ΔXu Z, Guo S, Tkachuk O, Nejati S, Razavi N, Argyros G (2024). "Cloud resource protection via automated security property reasoning". *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*. pp. 2170–2175.
342. ^ΔDisselkoe C, Kastner J, shaobo-he-aws, Hietala K, Wells A, Eline A, Moreno V, Palacios A, Markling M, Sze gheo N, yuan, Larsen MJ, Sharma S, B-Lorentz, Smith N, Vanderbleek S, Mamat A, Banchich A, Hakanson K, vasumv, Cecchetti S, Arakaki R, Flatt O, Meissl C, Bhakti, Rozek B, Garceda JV, TameIs J, Jones L. "cedar-policy/cedar". 2025. Available from: <https://github.com/cedar-policy/cedar>.
343. ^ΔBeyer D. Competition on software verification and witness validation: SV-COMP 2023. In: *Tools and Algorithms for the Construction and Analysis of Systems: 29th International Conference, TACAS 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2023, Paris, France, April 22–27, 2023, Proceedings, Part II*. Berlin, Heidelberg: Springer-Verlag; 2023. p. 495–522. ISBN: 978-3-031-30819-2.
344. ^ΔBeyer D, Petrenko AK. Linux driver verification. In: *Proceedings of the 5th International Conference on Leveraging Applications of Formal Methods, Verification and Validation: Applications and Case Studies - Volume Part II*. Berlin, Heidelberg: Springer-Verlag; 2012. p. 1–6. ISBN 9783642340314.

345. [△]Huang L, Ebersold S, Kogtenkov A, Meyer B, Liu Y (2023). "Lessons from formally verified deployed software systems (extended version)". *arXiv preprint arXiv:2301.02206*. [arXiv:2301.02206](https://arxiv.org/abs/2301.02206).
346. [△]Loughridge C, Sun Q, Ahrenbach S, Cassano F, Sun C, Sheng Y, Mudide A, Misu MRH, Amin N, Tegmark M (2024). "DafnyBench: A Benchmark for Formal Software Verification". *arXiv preprint arXiv:2406.08467*. [arXiv:2406.08467](https://arxiv.org/abs/2406.08467).
347. [△]Lohn E, Welleck S (2024). "miniCodeProps: a Minimal Benchmark for Proving Code Properties". *arXiv preprint arXiv:2406.11915*. Available from: <https://arxiv.org/abs/2406.11915>.
348. [△]Poesia G, Loughridge C, Amin N (2024). "dafny-annotator: AI-Assisted Verification of Dafny Programs". *arXiv preprint arXiv:2411.15143*.
349. [△]Li YC, Zetzsche S, Somayyaajula S (2025). "Dafny as Verification-Aware Intermediate Language for Code Generation". *arXiv preprint arXiv:2501.06283*. Available from: <https://arxiv.org/abs/2501.06283>.
350. [△]Misu MRH, Lopes CV, Ma I, Noble J (2024). "Towards ai-assisted synthesis of verified dafny methods". *Proceedings of the ACM on Software Engineering*. 1 (FSE): 812–835.
351. [△]Yang C, Li X, Misu MRH, Yao J, Cui W, Gong Y, Hawblitzel C, Lahiri S, Lorch JR, Lu S, et al. AutoVerus: Automated proof generation for Rust code. *arXiv preprint arXiv:2409.13082*. 2024.
352. [△]Aggarwal P, Parno B, Welleck S (2024). "AlphaVerus: Bootstrapping formally verified code generation through self-improving translation and treefinement". *arXiv preprint arXiv:2412.06176*.
353. [△]Silva Á, Mendes A, Ferreira JF (2024). "Leveraging Large Language Models to Boost Dafny's Developers Productivity". *arXiv preprint arXiv:2401.00963*. [arXiv:2401.00963](https://arxiv.org/abs/2401.00963).
354. [△]Song P, Yang K, Anandkumar A (2024). "Towards large language models as copilots for theorem proving in lean". *arXiv preprint arXiv:2404.12534*.
355. [△]Welleck S, Saha R (2023). "LLMSTEP: LLM proofstep suggestions in Lean". *arXiv preprint arXiv:2310.18457*. [arXiv:2310.18457](https://arxiv.org/abs/2310.18457).
356. [△]Li Z, Sun J, Murphy L, Su Q, Li Z, Zhang X, Yang K, Si X (2024). "A survey on deep learning for theorem proving". *arXiv preprint arXiv:2404.09939*. [arXiv:2404.09939](https://arxiv.org/abs/2404.09939).

Declarations

Funding: A. Gu is supported by the National Science Foundation (NSF) Graduate Research Fellowship under Grant No. 2141064. N. Jain is supported by NSF grants CCF:1900968, CCF:1908870, and by SKY Lab industrial sponsors and affiliates. A. Solar-Lezama is supported by the National Science Foundation

(NSF) and Intel Corporation through NSF Grant CCF:2217064. D. Yang is supported by the ONR YIP Award N000142412532.

Potential competing interests: No potential competing interests to declare.