

SoK: Towards Effective Automated Vulnerability Repair

Ying Li¹, Faysal Hossain Shezan², Bomin Wei¹, Gang Wang³, and Yuan Tian¹

¹University of California, Los Angeles

²University of Texas at Arlington

³University of Illinois Urbana-Champaign

{yinglee,davidwei23,yuant}@ucla.edu, faysal.shezan@uta.edu, gangw@illinois.edu

Abstract

The increasing prevalence of software vulnerabilities necessitates automated vulnerability repair (AVR) techniques. This Systematization of Knowledge (SoK) provides a comprehensive overview of the AVR landscape, encompassing both synthetic and real-world vulnerabilities. Through a systematic literature review and quantitative benchmarking across diverse datasets, methods, and strategies, we establish a taxonomy of existing AVR methodologies, categorizing them into template-guided, search-based, constraint-based, and learning-driven approaches. We evaluate the strengths and limitations of these approaches, highlighting common challenges and practical implications. Our comprehensive analysis of existing AVR methods reveals a diverse landscape with no single “best” approach. Learning-based methods excel in specific scenarios but lack complete program understanding, and both learning and non-learning methods face challenges with complex vulnerabilities. Additionally, we identify emerging trends and propose future research directions to advance the field of AVR. This SoK serves as a valuable resource for researchers and practitioners, offering a structured understanding of the current state-of-the-art and guiding future research and development in this critical domain.

1 Introduction

The relentless increase in software vulnerabilities poses a critical challenge for organizations, with consequences ranging from financial loss to reputational damage [1, 2]. This challenge is compounded by the limitations of manual repair processes, which are often slow, error-prone, and struggle to keep up with the volume of necessary fixes. Furthermore, numerous known vulnerabilities remain unpatched for extended periods, leaving systems exposed [3–5]. Automatic Vulnerability Repair (AVR) has emerged as a vital field to address these challenges, offering the potential to significantly reduce the time and resources to mitigate security risks.

AVR has witnessed significant progress, evolving from the template-based approach that applied predefined repairs

(often inflexible), random mutation, to more sophisticated approaches [6–8]. For example, constraint-solving techniques, often deriving constraints from symbolic execution, can be computationally expensive [9, 10]. Deep learning offers an end-to-end translation approach, but its effectiveness hinges on high-quality vulnerability repair datasets [11–13]. Currently, large language models (LLMs) show promise due to their vast training data, although limitations persist in function-level repairs and comprehensive program understanding [14–16].

Given the rapid growth of this field, a SoK is crucial to understand the landscape. This SoK aims to: (1) Bridge knowledge gaps by analyzing AVR approaches across synthetic data and real-world vulnerabilities. (2) Offer critical insights into the strengths and limitations of current work. (3) Propose promising future research directions. In this paper, we establish a taxonomy of existing AVR methodologies, grounded in fundamental principles, and conduct a quantitative comparison of current approaches. This SoK provides a comparative analysis of security patch generation methods, complexity levels, impact factors for successful repair, and vulnerability types, revealing a multifaceted landscape. We assess both synthetic and real-world vulnerabilities, finding no definitive “best” method; each has inherent limitations. For example, learning-based methods lack whole-program understanding, while non-learning methods struggle with precise constraint extraction, e.g., loop invariant synthesis. Based on our taxonomy and benchmarks, we analyze current research advancements, theoretical challenges, and future directions in AVR. Our findings indicate that performance across existing works varies significantly depending on the benchmark used, which often suffer from limited data points. We propose several compelling research directions to advance AVR, including the integration of hybrid approaches, interpretability of AVR, and the generation of high-quality specifications.

This SoK benefits both experts and practitioners. For security engineers, it provides (1) a systematic taxonomy to understand existing research; (2) analysis for leading repair methods; and (3) insights into limitations and future potential. Practitioners gain (1) a concise problem definition of vulnera-

bility repair and (2) a comprehensive evaluation with practical implementations of key approaches. In summary, we have the following key contributions:

- We define the core problem of security vulnerability repair and provide a comprehensive taxonomy and comparative analysis of repair approaches (Section 3), with a focus on patch generation techniques (Section 4). This discussion culminates in actionable takeaways and clearly identified open research questions.
- Our taxonomy classifies security patch generation approaches by 4 categories, 5+ strategies, 11+ methodologies, 5+ programming languages, and 12+ vulnerability types.
- We conduct quantitative evaluations of existing methods across C/C++ and Java benchmarks. We synthesize practical implications and highlight prevalent challenges in AVR, offering a critical assessment of the field’s progress.
- We explore the latest research trends, identify ongoing challenges and propose directions for future work. Through quantitative and qualitative analysis, we highlight gaps in current techniques and extract key findings to guide future advancements in AVR. We will keep our website (<http://sok-avr.github.io/>) updated with all the latest studies and findings to promote the research in this field.

Related Work. This study mainly differs from existing surveys in the following three aspects: (i) Existing works [17–23] focus on program repair from a general perspective rather than a security-specific one. For instance, Huang et al. [22] briefly discusses security vulnerabilities in the context of learning-based methods, Monperrus et al. [19] include only a subset of AVR-related papers, and Pinconschi et al. [20] demonstrated that general program repair tools perform poorly on vulnerability repair. These gaps highlight the need for our security-focused analysis. (ii) existing surveys on vulnerability repair focuses on LLM for AVR [14, 24], however, traditional methods (e.g., constraint-based methods) for AVR continue to evolve and demonstrate effectiveness in recent research [25, 26]; (iii) Prior surveys lack a systematic comparison across different AVR approaches. Our study bridges these gaps by providing both quantitative benchmarking across diverse datasets and qualitative analysis of various approaches, enabling researchers to better understand the current landscape and identify promising future research directions in AVR.

2 Preliminaries and Problem Setup

Figure 1 shows the timeline from vulnerability discovery to CVE publication and vulnerability repair process. The timeline from vulnerability discovery to CVE publication typically involves several key stages. Initially, the exact discovery time of a vulnerability is often unknown. CVEs are usually published after a vendor releases a patch, but in urgent cases or due to lack of cooperation, they may be disclosed before a patch is available. A vulnerability is classified as a

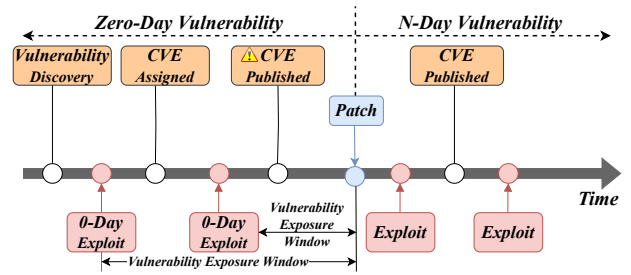


Figure 1: The timeline of vulnerability discovery, patch release, and exploit publication.

zero-day until a patch is released, after which it becomes an N-day vulnerability. The main goal of AVR is to minimize the vulnerability exposure window. The vulnerability repair workflow includes three key components: *vulnerability localization* (VL), *security patch generation* (SPG), and *security patch validation* (SPV). Below we detail each repair stage and further illustrate its relationship to program repair.

2.1 Vulnerability Localization

VL identifies the precise location or source of vulnerabilities within an application. This process involves pinpointing specific program points (i.e., specific locations within a program, e.g., lines of code, statements) where vulnerabilities exist and can be exploited, and then finding the point at which the “root cause” of the bug can be fixed [27].

Definition 1. We denote P as an original vulnerable program. Given V as the set of identified vulnerabilities within P , each vulnerability $v \in V$ is associated with a subset of program points $PT_v \subseteq PT$, where PT is the set of all program points in P . The vulnerability location L is the exact program point where the vulnerability v manifests can be fixed.

Accurately locating vulnerabilities is crucial for AVR, as crash points often differ from actual vulnerability locations, which are actionable for patching [27]. Early methods used slice-based VL: (1) static slicing analyzes control and data flow dependencies through dependency graphs but often includes irrelevant statements [28]; (2) dynamic slicing [29] records execution-specific data and control flow to extract relevant statements for given inputs [30]. Further, program state-based methods (e.g., delta debugging [31]) analyze runtime state changes by comparing states from successful and failed executions to pinpoint root causes. However, these methods are computationally expensive. In contrast, spectrum-based methods efficiently locate faults using only test coverage information by analyzing statement execution patterns in failed/passed tests and calculating suspiciousness scores [32].

Statistical reasoning like ConcFuzz [27], combines fuzzing near exploit paths with statistical analysis for precise VL. Building on this, learning-based methods further improve detection. For example, Xu et al. [32] locates vulnerabilities using reinforcement learning-guided fuzzing to generate coun-

terexamples, combined with spectrum-based methods for VL. Recently, LLMs have shown potential in VL, Zhang et al. [33] transforms the VL task into a sequence labeling problem, where each line of code is classified as either “vulnerable” or “non-vulnerable”, which has achieved 82.0% F1 score.

2.2 Security Patch Generation

A security patch is created to fix existing vulnerabilities, with the goal of making the program secure after applying it [34]. In contrast, non-security patches address issues like functionality bugs, feature updates or performance improvements. SPG is the process of creating critical patches for vulnerable code segments to effectively repair existing vulnerabilities.

Definition 2. Given a vulnerable program P , with identified vulnerabilities V at the specific locations L , a security patch P' is a modified version of P . The process of generating P' involves modifying P at each location l in L to fix the vulnerabilities present at that location.

2.3 Security Patch Validation

After generating security patches, the pivotal step is to verify if the patch fixes the given vulnerability without introducing any new bugs [35, 36] while maintaining correct program behavior [37, 38]. We define SPV as:

Definition 3. $SPV(P, P')$ evaluates if P' meets the following criteria: (i) fix the given vulnerability in P ; (ii) ensure no new bugs or vulnerabilities are introduced; (iii) preserve the original functionality of P . If all these three conditions are satisfied, the patch is deemed to be an effective security fix.

Li et al. [3] found approximately 12% of security patches are defective - with 7% fail to fully fix the target vulnerability, while 5% introduce new ones. The validation employs three main approaches: (1) Static analysis using code features [39, 40], expert rules [41], etc. to assess patch correctness without execution. For example, S3 [40] prioritizes patches by analyzing six static features measuring syntactic and semantic differences from buggy code; (2) Dynamic testing through test cases/exploits [42–45], execution traces [46] to test the patched code for verifying expected behavior; and (3) Formal methods [47, 48] like PATCHVERIF [47] used symbolic execution and physical invariant checks verify whether patch-introduced behavioral modifications meet expectations. Nowadays, LLMs like LLM4PatchCorrect [49] have been applied to predict patch correctness using contextual information from execution traces, bug descriptions. Despite this, human effort is still required in some AVR works (See in Section 3).

2.4 Security Vulnerability Repair

Problem Definition. Security Vulnerability Repair (SVR) involves generating P' by modifying a vulnerable program P at identified locations L to fix vulnerabilities V , ensuring that P'

satisfies the correctness criteria C . The objective is to find an appropriate repair function F that meets these requirements and $SPV(P, P') == True$.

The three key stages of SVR (i.e., VL, SPG, and SPV) work in tandem to ensure effective SVR. VL identifies the exact program points L of vulnerabilities V , providing inputs for SPG to generate patches that transform the vulnerable program P into repaired program P' . SPV verifies that P' fixes the vulnerabilities, preserves functionality, and introduces no new issues. These stages are interdependent, as errors in localization can propagate through SPG, while SPV feedback guides earlier steps, ensuring an iterative and reliable SVR.

2.5 Automatic Vulnerability Repair v.s. Automatic Program Repair

AVR is a subset of Automatic Program Repair (APR). While both aim to automate the repair process and share fundamental stages—localization, patch generation, and patch validation—they differ significantly in principles, assumptions, and input settings, particularly during the patch generation phase, which introduces unique challenges for AVR.

APR primarily targets general software issues, such as functional bugs [50, 51] and compilation errors [52–54]. In contrast, AVR focuses on vulnerabilities, requiring patches that both eliminate exploitable behaviors and preserve the program’s functional invariants. This dual requirement introduces additional complexity to the repair synthesis.

APR and AVR also differ in their assumptions. Traditional APR assumes program correctness can be fully captured through comprehensive test cases, while AVR relies on specific security properties and exploit conditions. Despite this difference, both approaches share the belief that effective repairs can be derived from patterns observed in historical fixes, enabling learning-based approaches in both fields.

The distinction in assumptions is most evident in their input requirements. APR typically relies on test suites to guide the repair process, while AVR operates with a single Proof-of-Concept (PoC) exploit that demonstrates the vulnerability. This limited information introduces unique technical challenges for AVR, as it must simultaneously ensure vulnerability elimination and functionality preservation without comprehensive test coverage.

3 Taxonomy of Automatic Vulnerability Repair Approaches

Selection Methodology. We conducted a systematic literature search focused on AVR using four search queries: “security patch generation”, “vulnerability repair”, “program repair + vulnerability” and “vulnerability patching”. The search was performed on [cspapers.org](https://www.cspapers.org) automatically, limiting results to the domains of Computer Security, Software Engineering, and Programming Languages. Initial results yielded 4,883 unique papers published between 2018 and 2024.9. To increase

Table 1: Properties and References of Security Patch Generation Approaches.

Category	Strategy	Methodology	Language ¹	Fix Level ²	Repair Input ³	Vulnerability Target ⁴	Test Suites ⁵	References	
Template Guided	Template Construction	Vulnerability Property	C/C++	S	FC	GN	✗	[55]	
			C/C++	S	S	Buffer Overflow	✗	[6, 56–58]	
			C/C++	S	P	Memory Error	✗	[7, 59–61]	
			C/C++	S	P	Integer Overflow	✗	[62, 63]	
			C/C++	S	S	Integer Overflow	✗	[8, 58]	
			Rust	S	P	Buffer Overflow	✗	[64]	
			C/C++	S	FC	Error Handling Bugs	✗	[65]	
			Solidity	BT	FC	Reentrancy, Acc Ctrl, Arithmetic, Uncheck LL Calls	✗	[66–68]	
			Solidity	S	FC	Arithmetic, Reentrancy	✗	[69]	
			Solidity	S	S, FC, P	Reentrancy, Miss.Input.Vald., Lock.Ether, Unhandl.Except.	✗	[70]	
			Java	S	S	SQL Injection	✗	[71]	
			Java	S	P	Null Pointer Deref.	✗	[72]	
		/	S	S	Denial of Service	✗	[73]		
		Summaried From Patches	Java	S	P	Cryptographic Misuses	✗	[74, 75]	
			Java	S	P	GN	✗	[76]	
			Java	S	S	GN	✗	[45]	
C/C++/Java	S		P	Null Pointer Deref.	✗	[25]			
Search Based	Mutation Based	Random Mutation	C/C++	S	S	GN	✓	[41, 78–81]	
			C/C++	S	P	GN	✓	[10]	
			Java	S	FL	Null Pointer Excep.	✓	[82]	
		Pattern Driven Mutation	C/C++	S	P	Information Flow Leakage	✗	[83]	
			Java	S	P	Null Pointer Excep.	✗	[84]	
			Solidity	S	P	GN	✗	[85]	
	Semantic Code Search	Security Patch Transplantation	PHP	S	P	Injection Vulnerabilities	✗	[86]	
			C/C++	S	S	GN	✗	[87, 88]	
			C/C++	S	P	GN	✓	[89]	
		Code Similarity	PHP	S	S	Access Control Bugs	✗	[90]	
			C/C++	B	FC	GN	✗	[91]	
			Java	S	S	Null Pointer Deref.	✗	[92]	
	Constraint Based	Constraint Extraction	Static Analysis	C/C++	S	FC	Resource/Memory Leak, Null Pointer Deref.	✗	[92]
				C/C++	S	FC	Memory Leak	✓	[93]
				C/C++	B	FC	GN	✓	[94]
				/	S	P	Null Pointer Deref., Data Leakage	✗	[95]
C/C++				S	P	Null Pointer Deref., Data Leakage	✗	[25]	
/				S	S	Denial of Service	✓	[96]	
Symbolic Execution			C/C++	S	FC	GN	✓	[97–99]	
			C/C++	S	P	GN	✓	[9]	
			C/C++	S	S	GN	✓	[100, 101]	
			C/C++	B	FC	GN	✗	[102]	
			Dynamic Analysis	C/C++	S	P	GN	✓	[9, 103]
				C/C++	S	P	GN	✓	[104]
C/C++				B	FC	GN	✗	[91]	
Formal Methods			C/C++	S	P	GN	✓	[10]	
			/	/	/	Trigger Action Programming	✓	[26]	
			Learning Based	Training	Deep Learning	C/C++	S	FC	GN
Java	S	S				GN	✗	[12]	
Fine-tuning	Adaption	C/C++		S	S	GN	✗	[106]	
		Interaction		C/C++	S	FC	GN	✗	[107]
		Zero-shot		C/C++/Java	S	FC	GN	✗	[14, 15, 108–112]
Prompt Engineering	Few-shot	C/C++		S	FC	GN	✗	[16, 113]	
		Chain-of-thought		C/C++/Java/Solidity	S	FC	GN	✗	[114–117]
LLM Integration	Multi-LLM	LLM-External Tool		C/C++/Java	S	FC	GN	✗	[118]
				C/C++/Java	S	FC	GN	✓	[117, 119]

¹ **Language:** Programming language that SPG method targets at, “/” indicates no specific language is targeted or not a programming language.
² **Fix Level:** The level at which security patches are generated: “S” for source code level, “B” for binary, and “BT” for bytecode.
³ **Repair Input:** Granularity of Input considered for SPG, from statements (S), functions (FC), and files (FL), to projects (P).
⁴ **Vulnerability Target:** Classifies methods by the vulnerabilities they address. “GN” means not target specific vulnerabilities.
⁵ **Test:** Whether test suites are required for SPG. “✓” indicates test suites are needed, “✗” means they are not.

relevance, we applied filtering criteria requiring papers to have a relevance score above 0.5 and contain at least one of the following keywords in their abstracts: “security”, “vulnerability”, “repair”, “patch”, or “fix”. This filtering process resulted in 267 unique papers. These papers underwent detailed review by two authors, spending about 45 minutes per paper to determine if the paper targets AVR, with weekly consensus meetings to resolve disagreements. Through this manual review process, we identified 23 core papers for AVR. Finally, we applied both

forward and backward snowballing like [120, 121] to examine referenced and citing papers, identifying additional 56 AVR papers, which cross 2006 to 2024.9, the process terminated when no new methodologies, strategies, programming languages on AVR were found in two consecutive iterations. The final collection comprises 79 papers, with 67 ranked A/A* on CORE2023 [122], a computing research venue ranking system where A/A* represents the top 20% venues.

Analysis Methodology. To analyze each paper, we em-

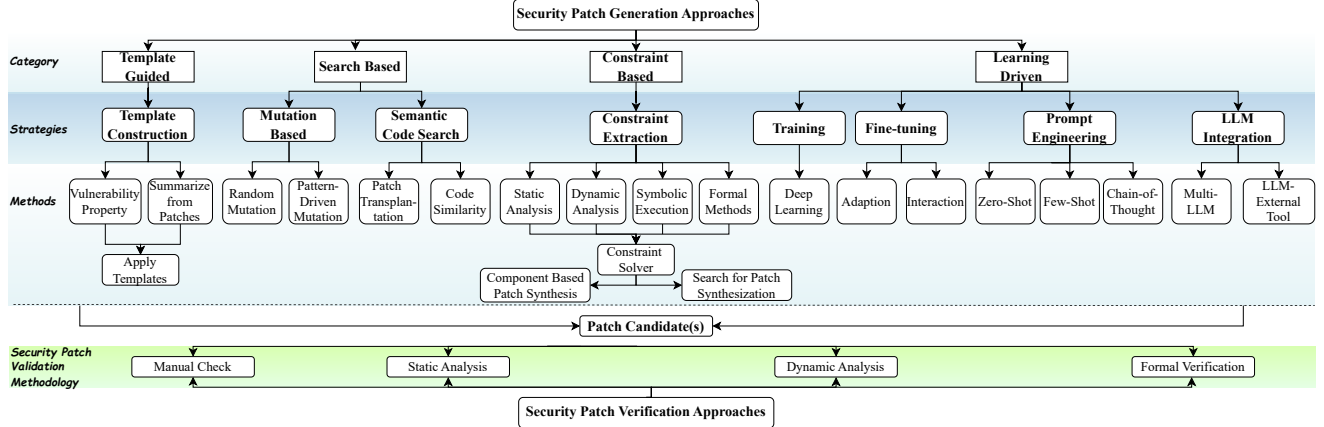


Figure 2: **Taxonomy** of Automated Vulnerability Repair approaches. Boxes with *blue* background present the taxonomy of security patch generation methods. Boxes with *green* background show the taxonomy of security patch validation approaches.

ployed a systematic approach to analyze each paper, examining three key dimensions: (1) The core components and methodologies used in AVR; (2) The effectiveness and target vulnerabilities types; (3) The strengths and limitations of each approach. We analyze each paper using a standardized rubric evaluating technical architecture, vulnerability types addressed, repair strategies with their theoretical foundations, experimental validation, and limitations. Uncertainty and inconsistencies were resolved through group discussion. This analysis revealed fundamental distinctions between learning-based and non-learning approaches. Within non-learning approaches, we identified three distinct categories: template-guided approaches using predefined templates, search-based approaches exploring security patch spaces or vulnerable states, and constraint-based approaches that formalize repairs through constraint extraction and solving. Further analysis within each category revealed distinct strategies, leading to our complete hierarchical taxonomy.

Figure 2 presents a hierarchical taxonomy, categorizing repair approaches into SPG and SPV (assuming successful vulnerability localization). SPG approaches are further classified into three levels: (i) broad categories (Template-Guided, Search-Based, Constraint-Based, and Learning-Driven), (ii) specific strategies within each category, and (iii) core methodologies. A detailed analysis of these approaches is provided in Section 4. SPV, classified into manual check, static analysis, dynamic analysis, and formal verification, shows a strong presence of manual in AVR practices (56 out of 79 AVR papers), either as a standalone approach or in combination with other methods. While VL is a critical prerequisite step for AVR, AVR papers often stress their core contribution to SPG, with 15 papers even assume known vulnerability locations.

Properties of Security Patch Generation approaches. From Table 1, we observe different features under different categories. For **Language**, 66 out of 79 papers focus on C/C++/Java, others focus on other languages (e.g., Rust). Notably, one paper [26] focus not on programming language but on natural

language. This language distribution aligns with real-world trends: C/C++ dominate memory-safety issues [123, 124], while Rust sees limited AVR coverage due to its newness. For **Fix Level**, only 4 provide AVR at the binary level, despite the importance of binary-level defense in deployed software systems [125]. Notably, 3 out of 6 non-learning-based studies on Solidity focus on bytecode-level instead of source-code level due to source code availability (0.3%) [68]. For **Repair Input**, current learning-based methods typically operate at statement-level or function-level repair input. This contrasts with other approaches, which offer support across all four input levels for repair, depending on the specifics of their implementation. Additionally, in the process of generating security patches, both learning-based and template-guided methods do not require the involvement of test cases or exploits. In contrast, other approaches may necessitate test cases/exploits, depending on their specific implementation details. For **Vulnerability Target**, 85.71% of the template-guided methods target one or more specific vulnerability types.

Various methods exhibit distinct strengths and weaknesses across different properties. For example, source-code level fixes are more straightforward but binary-level fixes save compilation time [77]. Function/statement-level repair input is faster than program-level but lacks global program context. Among 16 papers targeting specific vulnerabilities in C/C++, 62.50% focus on memory-related vulnerabilities such as memory leak and buffer overflow, which are crucial because exploiting them could give attackers access to the whole system [126]. However, few works focus on logic vulnerabilities, which are equally important, if not more so (e.g., DAO attack [127] results in \$8.5 million ETH loss). Detailed pros and cons of each method will be discussed in Section 4.

4 Security Patch Generation Approaches

In this section, we present typical SPG approaches: *Template-Guided*, *Search-Based*, *Constraint-Based*, and *Learning-*

Driven methods. We conclude each subsection by highlighting the *Takeaways* and *Open Problems* in this section.

4.1 Template-Guided

Template-Guided approaches for SPG rely on predefined repair templates or safety properties derived from human expertise and vulnerability characteristics. Their efficacy depends on the accuracy and coverage of the identified templates or properties, which determine the scope and precision of vulnerabilities that can be fixed. This section explores how various research efforts define and utilize these templates or properties to remediate security vulnerabilities.

A. Template Construction. There are two main directions for constructing the template:

A1. Templates based on Vulnerability Property. Different types of vulnerabilities have unique root causes and characteristics. Customizing security properties to address these facilitates the development of efficient repair templates. This approach ensures that each vulnerability is treated with a strategy that aligned to its nature. Initially, Weimer [55] proposed using an abstract behavioral model to generate security patches. Later, AutoPag [6] developed templates specifically for out-of-bound vulnerabilities, but it was only applicable when the vulnerability and the fix location were in the same function. Shaw et al. [56] focused on replacing unsafe APIs with safer alternatives to fix buffer overflows, but this template is effective only when the vulnerabilities are directly introduced by these unsafe APIs. Furthermore, BovInspector [57], Senx [58], and INTREPAIR [63] utilized safety properties of corresponding vulnerabilities related to specific vulnerabilities to derive patches, including buffer overflow, bad casts, integer overflows, and dangling pointers. Besides repairing statically, Exterminator [7] fixes memory bugs at runtime without requiring source code changes. Besides memory vulnerabilities, RegexScalpel [73] applies predefined repair templates (e.g., modifying quantifiers) to fix Regular Expression Denial of Service (ReDoS) vulnerabilities (e.g., nested quantifiers).

Templates derived from vulnerability properties can be represented or mined from various program analysis graphs (e.g., control flow graph, program dependency graph, abstract syntax tree (AST)) [128, 129]. These approaches primarily involve traversing the graph and modifying elements in fixed locations within the graph to achieve security patch generation. The basic idea has been used to fix common vulnerabilities, including memory management bugs (e.g., memory leak, double free, use after free [59, 61]), error handling bugs [65], null pointer dereferences [72], and integer overflows [8, 62]. Furthermore, SAVER [60] formulated the vulnerability fix the problem as a graph labelling problem to implement fixes.

Although the above researches have defined patterns to fix buffer overflow in C/C++/Java, the approaches cannot be directly used for fixing Rust buffer overflow vulnerabilities

due to its unique language features like ownerships and lifetimes [130]. Also, Rust vulnerabilities arise from interactions between its safe and unsafe sub-languages [131–133]. Rupa [64] defines the patterns to rectify the vulnerabilities by semantics-preserving program transformations. Similarly, the approaches are also not applicable on Solidity due to their domain-specific vulnerabilities, such as reentrancy attacks, gas limit issues, which are not common in other languages in traditional software. SGUARD [69] and CONTRACT-FIX [70] apply fixing templates to the source code to ensure the smart contract is free from those vulnerabilities, while EVMPatch [67], Smartshield [66], and Elysium [68] operates on bytecode level. However, they follow similar strategies, for example, SGUARD [69] add `nonReentrant` modifier for reentrancy vulnerabilities, EVMPatch [67] added checks to verify the caller’s address or permission on bytecode level.

A2. Templates Summarized from Patches. Different from templates guided by vulnerability properties, templates summarized from existing patches focus on generalizing successful strategies from real-world fixes. PAR [45], Vurle [76], Seader [75], and CONCH [25] summarized fix templates from patches for SPG. Ma et al. [74] summarized cryptographic misuses and corresponding fix patterns in Android. E9PATCH [77] defines four tactics for binary rewriting, proving effective in repairing vulnerabilities in binaries.

B. Template Application. Once constructed, the template is applied for patch generation. Note that not all templates are in source code format. When a template is not in source code, it must be mapped to concrete expressions (e.g., Senx [58]) to facilitate the generation of a patch in source code.

Takeaway I. *A template, whether enforcing security properties or serving directly as a repair template, often incorporates heuristic rules to address specific vulnerabilities or security property violations. This approach is effective to fix vulnerabilities that require minimal changes, have limited context dependency, and have clear root causes, such as switching to safer APIs or changing array sizes. While template-based methods lack flexibility, they are still a great starting point when applied to new contexts, such as binary-level repair [77]. Furthermore, templates can work with other repair methods throughout the AVR process, such as by extracting specific types of constraints [9] to enhance SPG.*

Open Problems I.

- **Template Mining:** *Most current template-guided AVR methods rely on hard templates (pre-defined and unchanging repair strategies)(e.g., [45, 60, 64, 66, 67, 69, 70]). While same type vulnerabilities often follow shared patterns, their specific manifestations can vary due to evolving contexts or exploitation techniques, necessitating flexible soft templates [45, 66, 76]. While existing research on patch generation explores using models to create repair templates, vulnerabilities often involve complex dependencies that require more flexible approaches. Thus, a soft template mining method is urgently needed.*
- **Security Property Inference:** *Security properties [58] can be*

viewed as templates to enforce stringent security conditions. Unlike typical fix templates, they guide other methods (e.g., solvers) to generate patches that meet complex requirements, such as sequence dependencies. Inferring these properties from various sources is essential for ensuring robust SPG.

4.2 Search Based

Heuristically search-based SPG is guided by defining heuristic rules to search and generate patches [134]. Two primary approaches within this category are *mutation-based* and *semantic code search*. Both establish a search space and select strategies accordingly. Mutation-based approach modify specific code sections for SPG. Conversely, the *semantic code search* seeks to identify code segments that are semantically similar but non-vulnerable code segments as potential solutions.

A. Mutation-Based Approach. This approach focus on modifying operations within a designated statement at a specified patch location, primarily by transforming the program’s AST. Mutations at the chosen repair site, involve alterations to operators, variables, APIs, and operation types within the AST. These mutations fall into two categories: *random mutation*, where changes are applied randomly, and *pattern-driven mutation*, guided by predefined patterns or rules.

A1. Random Mutation. Random mutation introduces a level of randomness during mutation, primarily using genetic programming and fuzzing. Genetic programming fundamentally operates through a cycle of *mutation* \Leftarrow *test*. This cycle continues until the mutated program successfully passes all test suites. GenProg [78–81] implements this concept, guided by the heuristic that “a program containing an error in one context likely implements correct behaviour elsewhere” [135]. This involves modifying the AST and recombining existing code snippets through mutation until all test suites are passed. This technique has demonstrated effectiveness in repairing vulnerabilities such as buffer overflow and format string issues [78–81]. However, the vast search space often leads to low-quality patches. To improve this, Tan et al. [41] suggested narrowing the search space by capturing disallowed modifications (i.e., modifications considered inappropriate or counterproductive for fixing bugs). However, obtaining “complete” anti-patterns is challenging. Anti-patterns are captured based on their frequency in bad patches and their rarity in correct ones. Consequently, patterns that are less prevalent in bad patches might be omitted.

Apart from genetic programming, fuzzing [136] is another technique for random mutation. VulnFix [10] employs snapshot fuzzing to mutate program states directly while inferring patch invariants. These invariants then guide the constraints on the patches, offering a more targeted approach for AVR.

A2. Pattern Driven Mutation. Pattern-driven mutation employs generalized patterns for modifications, such as adding null pointer checks, modifying boolean expressions, or altering

arithmetic operations. The mutation patterns address many issues but may lack proven effectiveness for specific cases. In contrast, template-guided repair uses templates that have been proven effective in other fixes. Sapfix [82] employs two mutation strategies involving statement insertion, specifically for `return null` or `null check` operations. Given that information leakage is closely tied to control flow [137], HyperGI [83] uses standard operators along with two novel ones like `!if` and `!for` to perform control flow-dependent mutations. However, these fixes are performed statically, Durieux et al. [84] explore the search space of possible patches for null pointer exceptions with metaprogramming [138], which analyzes the program’s real-time behavior and selects a repair pattern to mutate based on the current execution context. Beyond the realm of C/C++ and Java vulnerability repair, there is a growing interest in fixing vulnerabilities in Solidity, particularly for DeFi protocols. DeFinery [85] adopts three employs three operator mutation rules for this, notably without needing access to test cases or historical transaction data.

B. Semantic Code Search. When code exhibits similar vulnerabilities in comparable contexts or parallel functions, analogous patching strategies are often necessary. Since 20%-50% of large software systems consist of cloned code (i.e., code segments that are duplicated or highly similar across different parts of a software project) [139, 140], and patches are generally developed for the most recent version [141], earlier versions often have vulnerabilities that are contextually similar to these updates. Furthermore, software vulnerabilities frequently recur across different applications [142]. Consequently, by analyzing the characteristics of existing patches for known vulnerabilities or studying benign code that lacks these vulnerabilities, it becomes possible to repair similar unpatched vulnerabilities. Inspired by patch semantics, this approach has been applied in *Security Patch Transplantation* and *Security Patch Generation based on Similarity*.

B1. Security Patch Transplantation. Patch backporting is a key aspect of security patch transplantation [143], adapting the latest official patches for known vulnerabilities to older software versions, ensuring both compatibility [144, 145] and continued security [146]. Tools like SKYPORT [86], FIXMORPH [87], and TSBPORT [88] are developed for this purpose. SKYPORT targets web applications, whereas TSBPORT and FIXMORPH target Linux kernel vulnerabilities. These tools assume that official patches are semantically consistent across versions, but semantic differences can challenge automated patching, even when the exact fix location is known [88]. Another approach within security patch transplantation is to spread fixes across various implementations of the same protocol or functionality. For instance, PatchWeave [89] uses concolic execution to find and apply matching symbolic expressions in different software, enabling the transfer of fixes to similar issues in diverse environments.

B2. Security Patch Generation based on similarity. Unlike

patch transplantation, vulnerabilities often reappear across various software systems and functionalities [147]. Thus, historical fixes serve as valuable references for addressing similar vulnerabilities. The key idea is to identify and apply similar patches to fix analogous issues. FixMeUp [90] targets faulty access-control logic in web applications and inserts similar code where needed. Van et al. [92], match code fragments to ensure functional properties or explore similar code within the same program that meets necessary conditions. However, these repair methods require system interruption during patching. Later, Xu et al. [91] determine the similarity between the code to be fixed and the buggy code by comparing the semantics of the official patches with the vulnerable functions using program analysis and weakest precondition reasoning.

Takeaway II. *Mutation-based SPG uses random mutations that may not specifically tailored vulnerabilities but effective with clear patterns. Semantic code search uses related patches as references, but relies on robust matching. Similarly, similarity-based SPG explores large search spaces [148] by detecting code similarities, requiring sound verification (e.g., comprehensive test suites).*

Open Problems II.

- **Search Space Optimization.** *How can we automatically optimize the search space and adapt search strategies based on different scenarios (e.g., different types of vulnerabilities, different settings, i.e., across projects or within a project) [41, 78–81]?*
- **Sparse Vulnerability Patterns.** *How can we enhance the search methods for rare vulnerabilities [149], leveraging the unique properties of these vulnerabilities to guide program synthesis? Additionally, how can we leverage the unique properties of these rare vulnerabilities to guide the program synthesis effectively?*
- **Complex Logic and Structural Changes in Backporting.** *How can we effectively backport patches involving complex logic or structural changes, considering the need to understand intricate program semantics, handle significant alterations to algorithms or data structures, and dependencies or contexts that may not exist in the target version [86–89]?*

4.3 Constraint Based

Program constraints are limitations or conditions imposed on program behavior, inputs, outputs, or data structures [150]. In constraint-based SPG, the key lies in formalizing the constructed or extracted constraint that the synthesized patch is required to satisfy. The methods focus on formal constraint generation rather than constraint solving [151]. Constraint-based approaches could be summarized into two main steps: (i) Extract repair constraints. (ii) Employ constraint solvers or search for statements that fulfill the constraints. Our discussion is organized around various methods of the two steps.

A. Constraint Extraction. Constraint extraction involves analyzing the affected code to identify the specific requirements that any potential patch must adhere to. Constraint extraction methods can be classified into four approaches.

A1. Static Analysis Driven Constraint Extraction. Utilizing static analysis could identify constraints related to variables’ values, their interrelationships, and the function calls sequences. These identified elements can then be utilized to generate further constraints. MemFix [93] uses tpestate analysis [152] to track memory object states across branches, and generate patches for each state. It frames the problem as an NP-complete exact cover issue [153], focusing on patching memory vulnerabilities (e.g., double free). However, it is limited to modifying deallocation statements, excluding conditionals. CONCH [25] employs state propagation, utilizing CFG to encompass the full call chain for constraint extraction. This strategy aims to craft patches for Null Pointer Deferences (NPDs). Chida et al. [96] introduced a method by analyzing templates to maintain condition consistency with examples to generate constraints for repairing regex denial-of-service vulnerabilities. Symlog [95] employs Datalog-based [154] static analysis to generate structural constraints directly, effective across multiple languages due to Datalog’s declarative nature. However, the above methods only work when the application is open-source and focuses on source code. OSSPATCHER [94] assumes closed-source applications with available patches from open-source code. It employs a variability-aware AST to derive constraints, such as macro definitions and conditional compilations, and adapts these patches into binary form, aligning them with extracted constraints.

A2. Symbolic Execution Driven Constraints Extraction. Symbolic execution [155] executes the program with symbolic values, enabling the computation of constraints. Depending on the different objects of focus, different types of constraints can be computed through symbolic execution via tools such as KLEE [156]. The most representative method is to adopt symbolic execution to capture program semantic constraints that could pass all test cases [97, 98, 101]. Further, the constraint for memory writes and function calls [102], the weakest precondition, the extension of crash-free constraints [9] could also be computed by symbolic execution. To achieve better path coverage, concolic execution [157], which employs concrete input to drive symbolic execution, is also utilized to traverse the path driven by path constraints [99].

A3. Dynamic Analysis Driven Constraints Extraction. Dynamic analysis extracts runtime constraints by observing program behavior and potential vulnerabilities during execution. NOPOL [104] collects runtime data during test case execution to analyze the program behavior, identifying relevant variables and expressions, and encoded into a Satisfiability Modulo Theory (SMT) problem [158]. However, this method relies heavily on test suites quality, and most of the time, vulnerabilities may only have one exploit. To address this, ExtractFix [9] generates patches that generalize beyond a single exploit trace. Specifically, it adopts sanitizers [159, 160] to convert vulnerabilities into normal program crashes, allowing the extraction of precise condition constraints at the crash point.

A4. Formal Methods for Constraints Extraction. Formal methods use mathematical techniques and theories to verify and prove software correctness [161]. These methods involve constructing mathematical models that describe the behavior and properties of a program, allowing for the definition of precise constraints that the software should satisfy. VULMET [91] employs the weakest precondition reasoning to transform known patches into constraints. VulnFix [10] derives repair constraints via inductive inference [162]. VulnFix begins by collecting program state snapshots at vulnerability location during benign and vulnerable executions, then use Daikon [163] to infer initial invariants that distinguish benign states from vulnerable ones and CVC5 [164] verifies and refines these invariants to enhance the accuracy. Unlike approaches dependent on symbolic and concolic execution, these methods scale better to larger programs. Additionally, vulnerabilities not only exist in the code, but also stem from higher-level abstractions, like Trigger Action Programming (e.g., turn off lights when idle may conflict with a rule to keep them on), where logical inconsistencies or design flaws can introduce security risks. TAPFixer [26] detects rule conflicts via model checking [165] and resolves them using Negated-Property Reasoning, which analyzes violated properties to generate fixes such as adding delays, modifying conditions.

B. Patch Generation Based on Constraints. After extracting repair constraints, two typical methods for generating patches emerge. The first involves synthesizing a patch by finding a solution that ensures correct behavior and avoids the vulnerable state, typically using a component-based program synthesis approach. The second method involves searching the codebase for code that meets these constraints.

B1. Component-based Patch Synthesis involves selecting and combining predefined components (e.g., variables, operations) to automatically generate patches that satisfy expected behavior, typically using SMT solvers. The choice of SMT solvers depends on the constraint features and problem objectives. General-SMT solvers, like Z3 [166], are widely used in current research and applications, as evidenced by their use in a variety of works [9, 93–95, 97–99, 102, 104], in which [94] implement source-to-binary matching. For example, in `strcpy` function [167], a buffer overflow occurs if the `source` parameter is larger than `destination`. To prevent this, a constraint ensures `length(source) <= length(destination)`.

Besides checking the satisfiability of logical formulas over one or more theories, SAT [168] solvers handle Boolean logic satisfiability checks. Gopinath et al. [169] translate the constraint $pre - state \wedge code - constraints \wedge post - condition$ into boolean logic and solved by SAT solver.

However, when multiple patches meet the constraints using first-order logic, a second-order solver [170] enhances patch quality. For instance, EXTRACTFIX [9] employs pMaxSMT (partial MaxSAT) to handle both hard and soft constraints, while DirectFix [101] transforms the patch generation task

into a pMaxSAT problem, using patch complexity as a soft constraint. This approach helps produce simpler patches.

B2. Search Codebase For Patch Synthesization. Unlike *component-based patch synthesis*, this approach searches codebase for existing code that satisfies extracted constraints, bypassing the need to reassemble code components within the project [96]. More details are elaborated in Section 4.2.

Takeaway III. *Constraint-based SPG excels in capturing precise vulnerability requirements and program-wide properties, ensuring fixes satisfy specified invariants or constraints. Unlike heuristic methods, this approach is particularly effective for vulnerabilities requiring rigorous analysis of safety properties.*

Open Problems III.

- **Vulnerability-Specific Constraints.** *High-quality constraints are the key to constraint-based SPG (e.g., [94–96]), because it defines the properties that must hold true or false for repaired code. An open problem is how to accurately infer and differentiate between “benign” constraints and those specific to vulnerabilities.*
- **Security Specification Generation.** *Constraints can be considered a concretization of the specification, further refined to guide the implementation and operation of the system. However, specifications are often incomplete or underspecified [171], leading to ambiguity in the constraints to be enforced. Thus, an open challenge is to develop methods for inferring detailed specifications from existing code, textual information, etc. that can then be used to guide the constraint generation process.*

4.4 Learning Driven

Learning-driven SPG aims at utilizing learning methods to perform end-to-end security patch generation. Although in general bug repair, learning methods have been used to rank patches [172–175], extract templates for patch generation [176, 177], and end to end patch generation [94, 178–184], their application in AVR remains primarily limited to SPG.

Learning-driven SPG methods encompass *Training*, *Fine-tuning*, *Prompt Engineering*, and *LLM Integration*, with the latter three relying on large (code) language models. These models include encoder-decoder models that transform an input sequence $x = (x_0, x_1, \dots, x_T)$ into a corresponding output sequence; decoder-only models for next token prediction; and infilling models that generate contextual missing segments.

A. Training. Training involves building a model from scratch. In SPG, deep learning is particularly favoured for this purpose.

A1. Deep Learning. In deep learning, Neural Machine Translation (NMT) becomes fundamental for SPG. It utilizes encoder-decoder architectures, often using RNNs or transformers, learn mappings between source and target code sequences. For instance, Seqtrans [12] and VRepair [13] utilize NMT to learn rules from historical fixes (including vulnerability fixes and bug fixes) and apply them to future edits. Additionally,

Harer et al. [105] tackle vulnerability repair using adversarial learning with Generative Adversarial Networks [11]. They utilize an NMT model as the generator within the GAN framework, enabling training without paired examples of code.

B. Fine-tuning. Fine-tuning involves adjusting a pre-trained model to suit a specific task or domain. Since these models are not immediately suitable for AVR, significant customization and re-training are necessary to make them effective.

B1. Adaption. One method for fine-tuning is adaption, which focuses on adapting knowledge from one domain to another. Given the limited dataset for vulnerability fixes, the similarity between vulnerability fixes and bug fixing prompts the use of bug-fix datasets for pre-training [12, 13, 118, 185]. However, self-trained models on bug-fixing datasets often fail to capture rich programming features [50]. Some pre-trained models, trained on large codebases, can be fine-tuned to effectively address vulnerabilities. For example, VULRepair [106] fine-tuned pre-trained Programming Language models for SPG.

B2. Interaction. Fine-tuning can also involve learning through environmental interaction, i.e., reinforcement learning [186]. To ensure syntactic and semantic consistency, Islam et al. [107] used CodeBLEU [187] score and BERTScore [188] as rewards, guiding fine-tuning with Proximal policy optimization [189] to fine-tune CodeGen2 model [190] for SPG.

C. Prompt Engineering. Prompt engineering guides LLMs to produce optimal SPG outputs through designed inputs, encompassing *Zero-shot*, *Few-shot*, and *Chain-of-thought prompting*.

B1. Zero-shot prompting. Zero-shot prompting involves providing prompts without labeled data. A line of research [14, 15, 108–112] explores the capability of zero-shot AVR with different LLMs, including Codex [191], CodeT5 [192], GPT-4 [193], CodeGen2 [190] etc. Among these, Pearce et al. [15] combine error messages from static analysis tools in the prompts, NAVRepair [110] combines AST node-type information and error types in prompts, Liu et al. [111] combines CWE-ID, vulnerability location, root cause in prompts to guide SPG. Beyond software bugs, Ahmad [109] constructs prompts to fix hardware vulnerabilities.

B2. Few-shot prompting. Few-shot prompting uses a few labeled examples. VQM [113] provide example fixes of the CWE and data augmentation for repair. Though previously Fu et al. [16] tested gpt-3.5-turbo and gpt-4 with three repair examples per prompt and found that it failed to generate correct patches for all vulnerable functions in their dataset.

B3. Chain-of-thought prompting. Chain-of-thought (CoT) prompting improve LLMs’ logic-based task performance by simulating human reasoning, which address the limitations of other prompting methods in improving reasoning capabilities. Nong et al. [114], Khan et al. [115], VRpilot [117], and ContractTinker [116] utilized CoT to break down AVR process into different reasoning steps (e.g., identify vulnerability types, vulnerability causes analysis), all finally integrate

the results from previous reasoning results to guide SPG for these vulnerabilities. Also, VRpilot specify “the fix should not break any functionality of the function” in the prompt.

D. LLM Integration Approaches. LLM integration approaches enable collaboration between LLMs and other tools beyond pure prompting to enhance AVR performance.

D1. Multi-LLM Collaboration. To overcome the incompleteness in single model approaches, VulMaster [118] leverages two LLMs - CodeT5 as the fine-tunable backbone model and ChatGPT as the supplementary model for generating CWE examples and fixes, doubling the repair accuracy compared to approaches that rely solely on fine-tuned CodeT5.

D2. LLM-External Tool Collaboration. To enhance LLMs’ ability to understand vulnerabilities (e.g., counterexamples), they are integrated with external tools. ESBMC-AI [119] uses bounded model checker ESBMC [194] to locate vulnerabilities and generate counterexamples (including stack traces, line numbers, and variable names), then feed the counterexamples and original code into LLMs for SPG, and iteratively verifies and refines the fixes through ESBMC until the code passes the verification. VRPilot [117] integrates external tools output (e.g., code sanitizers) into the prompt to guide SPG.

Takeaway IV. *Learning-driven methods have transformed AVR by: (1) moving beyond traditional heuristic approaches and templates; (2) improving generalizability and offering more flexibility in inputs through fine-tuning, prompt engineering, and LLM integration approaches. Especially, prompts that provide contextual information (e.g., root causes) guide LLMs to generate more precise patches (e.g., [117]); (3) LLM integration approaches enable complementary capabilities (e.g., different models and tools) to work together for precise analysis and iterative SPG, enhancing robustness.*

Open Problems IV.

- **Component Dependency Analysis and Integration.** *Improving SPG requires analyzing complex program dependencies and their security implications ((e.g., [15, 110, 111, 117])). This includes: (1) understanding multi-component interactions like API calls and structure definitions, (2) tracking inter-procedural data flows, and (3) validating dependencies against security specifications. Thus, developing multimodal methods to analyze these may lead to more accurate, contextually informed repairs.*
- **Common Challenges in Learning-based Security Analysis.** *Common problems like data sparsity, input length limitations, and training data availability in learning-based security analysis are even more challenging in AVR. For example, while fine-tuning has proven effective in AVR, the lack of high-quality vulnerability repair datasets (e.g., [106])—comprising not only official patches but also semantically equivalent patches and clear evaluation metrics—remains a significant hurdle.*
- **Generated Security Patches Selection.** *In non-deterministic models like GPT-4, automatically validating and selecting the correct security patch from inconsistent outputs remains challenging. Unlike general bug fixes, security patches require*

Table 2: Evaluation results on C/C++ Vulnerability Benchmark. This table presents data in the format: Successful Repairs / Tests Passed / Successful Compilations / Total Tests Conducted (RSR). ‘NA’ means the tool is not applicable to the benchmark.

Benchmark	#SARD	#ExtractFix	#VulnLoc
Senx [58]	N/A	2/2/6/30(6.7%)	4/4/8/43(9.30%)
VulnFix [10]	N/A	13/13/19/30(43.33%)	20/20/28/43(46.51%)
ExtractFix [9]	N/A	21/21/25/30(70.00%)	16/16/18/43(37.21%)
VRPilot [117]	N/A	17/17/22/30(56.67%)	18/18/25/43(41.86%)
InCoder [205]	79/97/120/1000(7.90%)	0/2/5/30(00%)	0/3/7/43(00%)
Gemini-Pro [206]	121/153/971/1000(12.10%)	1/1/25/30(3.33%)	0/0/32/43(00%)
GPT-4-1106-preview [193]	966/982/1000/1000(96.6%)	14/14/20/30(46.67%)	15/15/23/43(34.88%)

Table 3: Evaluation results on Java Vulnerability Benchmark. **VJBench-Trans(R)**:renaming transformation, **S**: structure change, **R+S**: changed by both. Data Format: Successful Repairs / Tests Passed / Successful Compilations / Total Tests Conducted (RSR).

Benchmark	#VJBench	#VJBench-Trans(R)	#VJBench-Trans(S)	#VJBench-Trans(R+S)	#VUL4J
GPT-4-1106-preview [193]	5/5/20/23(21.74%)	5/5/22/23(21.74%)	5/7/20/23(21.74%)	3/5/20/23(13.04%)	26/30/56/70(37.14%)
VRPilot [117]	6/9/20/23(26.08%)	6/9/22/23(26.08%)	6/8/20/23(26.08%)	5/6/20/23(21.74%)	29/31/53/70(41.43%)
Gemini-Pro [206]	2/2/14/23(8.70%)	2/2/13/23(8.70%)	2/2/14/23(8.70%)	2/2/13/23(8.70%)	7/10/38/70(10.00%)
CodeT5 [192]	0/0/0/23(00%)	0/0/1/23(00%)	0/0/0/23(00%)	0/0/0/23(00%)	2/2/6/70(2.86%)
InCoder [205]	2/2/10/23(8.70%)	1/1/6/23(4.35%)	1/1/5/23(4.35%)	1/1/6/23(4.35%)	6/10/20/70(8.57%)
Fine-tuned-CodeT5 [14]	3/4/17/23(13.04%)	3/3/17/23(13.04%)	3/3/16/23(13.04%)	2/2/18/23(8.70%)	2/10/48/70(2.86%)
Fine-tuned-InCoder [14]	3/4/15/23(13.04%)	3/3/16/23(13.03%)	4/4/16/23(17.39%)	2/2/17/23(8.70%)	6/12/51/70(8.57%)

complex analysis of exploitability and system-wide security implications, making automated validation methods still insufficient. This validation process currently relies heavily on human expertise (e.g., [14, 115]).

- **Iterative AVR.** Tools feedback incorporation has proven AVR improvement (e.g., [117, 119]), how can we leverage LLM integration usages (e.g., LLM multi-agents) with long memories to enhance it not only limited to statement/function?

5 Benchmark Evaluation

In Sections 3 and 4, we explore AVR advances and evaluate SPG methods. Here, our goal is to understand their practical strengths and weaknesses.

5.1 Benchmark Evaluation Setup

Dataset Selection. We benchmark SPG methods by curating a dataset that spans various languages and weakness types. After reviewing vulnerabilities with fixes datasets ([9, 14, 27, 195–202]), we strategically select 4 datasets. Our selection criteria include: (1) prevalence in AVR research, (2) programming language diversity, (3) mix of synthetic and real-world vulnerabilities, and (4) real-world vulnerabilities test cases/exploits availability for verification. Finally, we selected: (1) synthetic dataset \mathcal{D}_{SARD} [203] (used in learning-based AVR evaluation [107, 114, 204]). For our evaluation, we randomly sampled 1,000 samples from this dataset. (2) real-world Java vulnerabilities \mathcal{D}_{Vul4J} (#79) [202], $\mathcal{D}_{VJBench}$ (#42) [14] and its transformations. (3) real-world C/C++ vulnerabilities $\mathcal{D}_{ExtractFix}$ (#30) [9], and $\mathcal{D}_{Vulnloc}$ (#43) [27], totaling 48 distinct vulnerabilities across $\mathcal{D}_{ExtractFix}$ and $\mathcal{D}_{Vulnloc}$. See details in Appendix A.

Inclusion & Exclusion Criteria. We evaluate contemporary AVR approaches, including but not limited to work from premier conferences. Tools were filtered based on artifact availability and must be: (1) be publicly available or accessible from authors, (2) executable, and (3) reproducible and applicable to selected dataset.

Ultimately, we evaluated 10 tools, including property-guided approach Senx [58], constraint approach ExtractFix [9], combined approach (search + constraint) VulnFix [10], and learning-based methods, including code infilling models InCoder [205] and CodeT5 [192], their fine-tuned versions [207], and generative models using prompt engineering, specifically GPT4-1106-preview [193] and Gemini-Pro [208], and VRPilot [117] which leverages CoT prompting with error messages, to ensure the fairness for evaluation, we also applied GPT4-1106-preview for VRPilot instead of GPT-3.5-turbo as claimed in their paper. The experimental settings, tools applicability, and evaluation steps please refer to Appendix B.

Evaluation Metrics. We evaluate the capability of SPG methods by measuring their repair success rate (RSR):

$$RSR = \frac{\text{Number of Vulnerabilities Successfully Repaired}}{\text{Total Number of Vulnerabilities to be Repaired}} \quad (1)$$

Note that for non-deterministic models, we generate 3 outputs for each input; as long as one of them passes the test, we include it in the calculation.

Evaluation Results. Tables 2 and 3 show the number of successful repairs, oracle (tests/exploits) passed, successful compilations, and total test counts. These metrics, along with RSR, highlight failure stages for each benchmark. We filtered Java datasets¹ due to learning methods’ limitations with multi-file modifications, retaining 23/42 vulnerabilities in

¹If the same modification is applied across different files, we still count it as a single modification.

$\mathcal{D}_{VJBench}$ and 70/79 vulnerabilities in \mathcal{D}_{Vul4J} . All real-world C/C++ benchmarks were kept since they could be tested on non-learning methods.

5.2 Highlighted Findings

As shown in Table 2 and Table 3, AVR performance varies significantly across benchmarks. GPT-4-1106-preview excelled on \mathcal{D}_{SARD} with an RSR of 96.6%. However, its effectiveness dropped below 50% in other benchmarks, highlighting its limitations with real-world vulnerabilities. Other learning-based models had RSRs under 20%. In contrast, VulnFix and ExtractFix handled real-world scenarios better, though their RSRs were not exceed 70%. Senx, limited to three vulnerability types, resulting in a low RSR on diverse benchmarks. Additionally, while VRPilot demonstrates improved RSR when integrated with external tool outputs, its performance remains below that of non-learning-based methods.

Learning-based methods show higher compilation rates than RSR but struggle with actual vulnerability fixes. For instance, on \mathcal{D}_{Vul4J} , GPT-4-1106-preview compiled 56 patches with 30 passing test cases but only 26 correctly fixed vulnerabilities, highlighting insufficient test coverage. Despite improvements in compilation (0 to 17 compilable patches in $\mathcal{D}_{VJBench}$) with fine-tuned models (e.g., fine-tuned CodeT5), successful repairs remains low (0 to 3).

Due to insufficient repaired samples in $\mathcal{D}_{VJBench}$ for meaningful comparisons, we applied 3 mutation strategies: variable swapping, condition reconstruction, and loop transformation (Appendix C) on \mathcal{D}_{SARD} with 200 random samples, still in high RSRs with GPT-4-1106-preview (97.5%, 96.15%, and 96.43% respectively).

Finding I. *SPG methods show varying effectiveness across benchmarks with no consistently dominant approach. The robustness of learning-driven method in our evaluation is excellent. Iterative LLM-external tool integration enhances AVR.*

Open RQ I. *How can we improve learning-based methods to maintain high compilation rates while ensuring functional consistency and enhancing security? Given commonly used RSR only measure the proportion of successful repairs, how can we develop comprehensive evaluation metrics that better reflect the real-world effectiveness (e.g., also consider the vulnerability severity)?*

Table 4: Successful repairs by Scope of Change and Code Dependencies on $\mathcal{D}_{ExtractFix}$ and $\mathcal{D}_{Vulnloc}$, **SH**: Single-hunk, **MH**: Multi-hunk, **Intra/Inter/Other**: corresponding dependencies

Tool	Scope of Change		Code Dependencies		
	SH	MH	Intra	Inter	Other
VRPilot	13/25	9/23	8/15	7/23	7/10
GPT-4-1106-prev	11/25	8/23	8/15	5/23	6/10
ExtractFix	15/25	6/23	8/15	10/23	3/10
VulnFix	12/25	8/23	6/15	10/23	4/10

To understand which changes make vulnerabilities more likely to be repaired, we examine the following dimensions: *Scope of Change*, and *Code Dependencies*. Since benchmark $\mathcal{D}_{ExtractFix}$ and $\mathcal{D}_{Vulnloc}$ feature a broad application of various methods, our analysis primarily focuses on it. *Scope of Change* refers to the number of modified code blocks, categorized into single-hunk (modifications confined to one contiguous code block) and multi-hunk changes (changes across multiple separated blocks). *Code Dependencies* are analyzed at three levels: intra-procedural dependency (fixed entirely within a single function or method, without needing to involve other components or functions.), inter-procedural dependency (modifications within the vulnerable function involve components outside of this function), and others (e.g., built-in features).

Table 4 indicates minimal impact of *scope of changes* on repair success across methods. VRPilot shows worse performance (RSR 30.43%) with inter-dependencies compared with ExtractFix and VulnFix.

Finding II. *“Scope of changes” is not a reliable metric to indicate the difficult level of AVR (although commonly used to assess the difficulty of repairing general bugs). In contrast, “code dependencies” have a more significant impact on RSR.*

Open RQ II. *How can we enhance the interpretability of AVR methods across different dimensions (e.g., vulnerability dependencies, statement length, modified tokens and lines, control structures, and vulnerability logic)?*

```

1 @@ -1632,6 +1632,13 @@ JPEGSetupEncode(TIFF* tif)
2     "Invalid horizontal/vertical sampling value←
3     ");
4     return (0);
5 }
6 + if( td->td_bitspersample > 16 )
7 + {
8 +     TIFFErrorExt(tif->tif_clientdata, module,
9 +                 "BitsPerSample %d not allowed ←
10 +                 for JPEG",
11 +                 td->td_bitspersample);
12 +     return (0);
13 + }

```

Listing 1: Patch for CVE-2017-7601, which was not repaired by learning-based methods but by VulnFix

We further investigated vulnerabilities that could be repaired by non-learning-based but not by learning-based approaches. A closer look at the vulnerabilities that learning-based methods failed to repair reveals complex issues, including intricate control flow conditions, program dependencies, non-obtainable structure members, global variables, or specifications. For example, in Listing 1, the vulnerability is caused by shift operations with excessively large exponents, which is repaired by a single-hunk modification. This stems from a lack of input validation for `td->td_bitspersample`, allowing values exceeding JPEG specification [209]’s 16-bit limit and causing integer overflow. VulnFix and ExtractFix can apply the correct repair by computing correct constraints, e.g., VulnFix could

generate the correct patch invariant at the crash point, which is further utilized for SPG, while learning-based methods miss details regarding the specification. However, if such specification is provided in prompts, GPT-4-1106-preview can correctly perform the repair. We also analyzed the cases learning-based methods repaired but non-learning-based methods did not, and cases that both approaches repaired; see details in Appendix D.

Finding III. Learning-based methods lack a comprehensive understanding of the entire program, struggling with vulnerabilities involving complex interrelations, atypical constructs, extensive program-wide constraints, or implicit constraints defined in specifications. In contrast, non-learning-based methods excel by leveraging the broader program context and computing critical constraints or invariants.

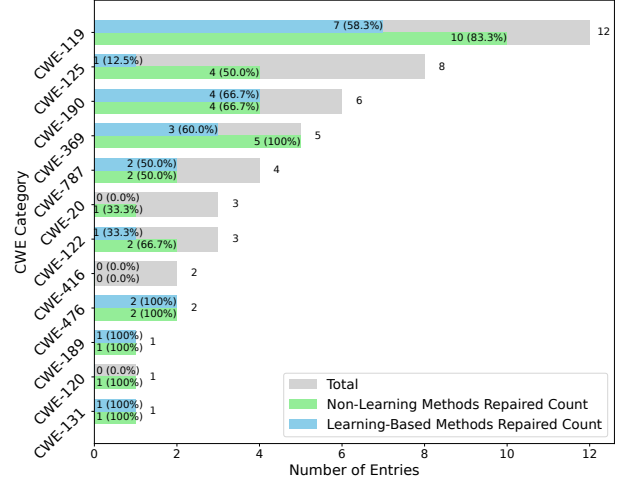
Open RQ III. How can we systematically identify and extract vulnerability root causes to generate precise specifications that guide learning-based repairs? Additionally, how can we integrate contextual dependencies and derive constraints from textual sources (e.g., specification documents) to enhance AVR?

We analyzed the weakness different methods successfully fixed, focusing on performance across CWE categories. For D_{SARD} , GPT-4-1106-preview has achieved an RSR 100% in 52/60 CWEs, likely due to the simple contextual information.


Figure 3 shows the repair distribution using learning and non-learning based approaches of CWE types across all real-world benchmarks. As shown in Figure 3(a), on each CWE type, no learning-based method performs better for any CWE. Specifically, both achieve 100% RSR on CWE-476 (NPDs) and CWE-189 (Numeric Errors). For example, in the case of CWE-189, like CVE-2016-10094, a numeric error was fixed by modifying the comparison from ≥ 4 to >4 in the code, preventing heap-based overflow. For CWE-476, the main fix involved adding the necessary null check to prevent dereferencing null pointers. Figure 3(b) shows that on each CWE with 2+ samples, the RSR is not greater than 75%. Note that although some CWEs show high RSR, the limited number of data points makes it unreliable to draw any strong conclusions.

Moreover, benchmarks across different programming languages focus on distinct CWE types, C/C++ benchmarks are primarily concerned with memory-related, like CWE-119 (Improper Restriction of Operations within Memory Buffer Bounds), while for Java, due to its garbage collections mechanism, tends to focus more on application-layer vulnerabilities (e.g., CWE-611, Improper Restriction of XXE Reference). The diverse application-layer vulnerabilities make it hard to design non-learning AVR tools for multiple CWEs in Java.

Finding IV. Learning-based AVR methods lag behind traditional approaches overall, though they excel at specific CWEs. The performance gap between real and synthetic datasets suggests synthetic data may inadequately evaluate AVR methods.



(a) Repair on Real world C/C++ Benchmarks



(b) Repair on Real world Java Benchmarks

Figure 3: Total vs. (Non-)Learning-Based Repair Counts Across CWE Categories in real-world benchmarks.

Open RQ IV. Can we establish a comprehensive benchmark for AVR? While some CWEs show high RSR, limited data makes strong conclusions unreliable.

6 Future Directions

Based on our literature review (Section 4) and evaluation (Section 5), we discuss future research directions here.

D1. Hybrid Approaches for SPG. Based on Section 4.4 and Section 5.2, A hybrid approach that combining fuzzing, program analysis, etc. with LLMs for SPG may presents a promising direction. For example, while program analysis can identify critical paths and dependencies, it may struggle with scalability in large codebases. Here, LLMs can optimize the

search space. This synergy between different techniques can be integrated into repair process. Also, fine-tuning LLMs with insights from program analysis may also helpful for context-aware and precise repair. As discussed in Section 5.2, artifacts like documentation and bug reports help with in-context learning. Additionally, LLM could generate constraints or patch invariants based on the analysis of vulnerable and benign states. The effectiveness of LLMs for invariant generation has been demonstrated in recent works [210, 211] but not tried in AVR yet. Compared to other software security tasks, hybrid approaches for AVR are more challenging as they must generate patches that account for dependencies, maintain semantic consistency, and avoid introducing new vulnerabilities.

D2. Domain-Specific AVR. As discussed in Section 3, while logic vulnerabilities receive less attention, they can be more severe, as privacy handling flaws may lead to data breaches and significant losses. General-purpose AVR tools often struggle with logic vulnerabilities due to their intricate logic, which may involve subtle design flaws or complex program behavior that isn't as straightforward to detect or repair as memory-related issues. Developing domain-specific methods for logic vulnerability repair, incorporating domain-specific languages, formal specification, and security properties and multi-modality information, is the foundation for reliable patches in domain-specific AVR (e.g., TAPFixer [26]).

D3. AVR Benchmarks. While numerous datasets on vulnerabilities and their fixes exist (Table 5), most lack context and exploits. Current AVR datasets (Section 5.2) have limited CWE coverage. AVR needs comprehensive benchmarks with vulnerable code, patches, working exploits, and reproducible test environments, challenging to create and maintain.

D4. Interpretability of AVR. As discussed in Section 5.2, current AVR methods, especially learning-based ones, lack interpretability regarding which vulnerabilities are easier/harder to repair. Simple metrics like changed hunks or lines provide limited insights. More comprehensive metrics (control flow complexity, data flow changes, statement length) are necessary to improve interpretability of AVR performance.

D5. Automatic Generation of High-quality Specifications. As discussed in Section 4, specifications are essential for extracting security properties and constraints. Section 5.2 further demonstrates that integrating more detailed information about vulnerable program points can enhance prompt engineering in learning-based methods, thereby increasing RSR. These insights highlight the need for better security-related specifications, as they define what constitutes a secure state and are crucial for enhancing the effectiveness of AVR. For example, in software documentation, security properties might be spread across different sections, using LLMs/NLP and automated reasoning techniques to extract and find potential violations may lead to security specific specifications.

D6. Verifier for Generated Security Patches. As discussed in Section 4, existing approaches still need human intervention in

SPV. When the search space is large (e.g., LLMs can generate multiple patches) potentially including a correct one. Also, incorrect patches can turn non-exploitable bugs into severe vulnerabilities[212] (e.g., the commit[213] fixed a memory leak but introduced a new double-free vulnerability). To this end, we advocate for further developments in patch verification methods to improve AVR efficiency, which may related to formal verification techniques.

D7. LLM In AVR. With the advancements in LLMs, LLMs might be used in any stage (VL, SPG, SPV) of AVR. As discussed in Section 2 and Section 4, currently, the use of LLMs in these areas is mostly limited to analyzing the code itself or incorporating a small amount of vulnerability-specific information like vulnerability types. However, since LLMs are prone to hallucination which may lead to inaccurate outputs, LLMs for AVR cannot solely rely on the LLMs' capabilities themselves. It is crucial to integrate them with traditional complementary techniques. **For VL**, LLMs can help filter relevant statements from dependency graphs in slice-based VL (discussed in Section 2), reducing noise and improving the precision of the analysis. Moreover, LLMs can assist counterexample generation by generating inputs that expose differences in predicates that may correlate with crashes. These counterexamples help rank predicates based on their likelihood of being root causes, improving the precision of localization. **For SPG**, as discussed in D1, LLM can be integrated with other techniques and further for patch synthesis in SPG. **For SPV**, though currently LLMs have been used for self-correction or self-checking [214], these approaches alone cannot prove correctness. However, LLMs can serve multiple roles in enhancing verification approaches. First, LLMs can extract semantic properties and security invariants from patches, which can then be formally encoded as verification conditions for automated theorem provers or model checkers. Second, LLMs can assist in translating patches into intermediate verification languages or logical formulas that are more amenable to automated reasoning. Third, when verification fails, LLMs can help analyze counterexamples generated by theorem provers and suggest refinements to the verification conditions or patches. This combination leverages LLMs' natural language understanding and code analysis capabilities while relying on the mathematical rigor of automated reasoning for formal verification, making SPV more robust. In whole AVR process, LLM agents serve to decompose complex repair tasks into actionable steps, interact with vulnerability analysis, and maintain critical context including code structure, repair history, etc. to facilitate effective AVR. Please refer to Appendix E for more future directions.

7 Conclusion

We present an SoK concerning SPG in AVR, including a comprehensive taxonomy of their characteristics and trade-offs, with future research directions theoretically and empirically.

Acknowledgements

We extended our sincere gratitude to our shepherd and the anonymous reviewers for their invaluable feedback and constructive suggestions. This work is supported in part by the National Science Foundation (NSF) Awards 2323105, 2317184, 2411153, 2229876, 2055233, and Cisco Research. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of sponsors.

Ethics Considerations

Our study focuses on the analysis and of existing vulnerability repair techniques. It does not involve experiments or interactions that could raise ethical concerns. All data used in this research is derived from publicly available sources, and no new vulnerabilities are introduced or exploited as part of this work.

Compliance with Open Science Policy

Our paper fully adheres to the open science policy introduced by USENIX Security. The used tools and datasets can be found at <https://zenodo.org/records/14750182>.

References

- [1] Forbes Technology Council. Zero-day vulnerabilities: 17 consequences and complications. <https://www.forbes.com/sites/forbestechcouncil/2023/05/26/zero-day-vulnerabilities-17-consequences-and-complications>. Accessed on: Jan, 2024.
- [2] Yonghee Shin, Andrew Meneely, Laurie Williams, and Jason A Osborne. Evaluating complexity, code churn, and developer activity metrics as indicators of software vulnerabilities. *IEEE TSE*, 2010.
- [3] Frank Li and Vern Paxson. A large-scale empirical study of security patches. In *Proc. ACM CCS*, 2017.
- [4] Xiaodan Li, Xiaolin Chang, John A Board, and Kishor S Trivedi. A novel approach for software vulnerability classification. In *Proc. IEEE RAMS*, 2017.
- [5] Makenzie Bunning. Vulnerability remediation timelines: 7 best practices. <https://www.ninjaone.com/blog/vulnerability-remediation-timelines-best-practices/>, 2024. Accessed on: Mar, 2024.
- [6] Zhiqiang Lin, Xuxian Jiang, Dongyan Xu, Bing Mao, and Li Xie. Autopag: towards automated software patch generation with source code root cause identification and repair. In *Proc. ACM CCS*, 2007.
- [7] Gene Novark, Emery D Berger, and Benjamin G Zorn. Exterminator: automatically correcting memory errors with high probability. In *Proc. ACM PLDI*, 2007.
- [8] Zack Coker and Munawar Hafiz. Program transformations to fix c integers. In *Proc. ACM/IEEE ICSE*, 2013.
- [9] Xiang Gao, Bo Wang, Gregory J Duck, Ruyi Ji, Yingfei Xiong, and Abhik Roychoudhury. Beyond tests: Program vulnerability repair via crash constraint extraction. *ACM TOSEM*, 2021.
- [10] Yuntong Zhang, Xiang Gao, Gregory J Duck, and Abhik Roychoudhury. Program vulnerability repair via inductive inference. In *Proc. ACM ISSTA*, 2022.
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 2020.
- [12] Jianlei Chi, Yu Qu, Ting Liu, Qinghua Zheng, and Heng Yin. Seqtrans: automatic vulnerability fix via sequence to sequence learning. *IEEE TSE*, 2022.
- [13] Zimin Chen, Steve Kommrusch, and Martin Monperrus. Neural transfer learning for repairing security vulnerabilities in c code. *IEEE TSE*, 2022.
- [14] Yi Wu, Nan Jiang, Hung Viet Pham, Thibaud Lutellier, Jordan Davis, Lin Tan, Petr Babkin, and Sameena Shah. How effective are neural networks for fixing security vulnerabilities. In *Proc. ACM ISSTA*, 2023.
- [15] Hammond Pearce, Benjamin Tan, Baleegh Ahmad, Ramesh Karri, and Brendan Dolan-Gavitt. Examining zero-shot vulnerability repair with large language models. In *Proc. IEEE S&P*, 2023.
- [16] Michael Fu, Chakkrit Kla Tantithamthavorn, Van Nguyen, and Trung Le. Chatgpt for vulnerability detection, classification, and repair: How far are we? In *IEEE APSEC*, 2023.
- [17] Martin Monperrus. Automatic software repair: A bibliography. *ACM Computing Surveys (CSUR)*, 2018.
- [18] Luca Gazzola, Daniela Micucci, and Leonardo Mariani. Automatic software repair: A survey. In *Proc. ACM/IEEE ICSE*, 2018.
- [19] Martin Monperrus. *The living review on automated program repair*. PhD thesis, HAL Archives Ouvertes, 2018.
- [20] Eduard Pinconschi, Rui Abreu, and Pedro Adão. A comparative study of automatic program repair techniques for security vulnerabilities. In *Proc. IEEE ISSRE*, 2021.
- [21] Emily Winter, Vesna Nowack, David Bowes, Steve Counsell, Tracy Hall, Sæmundur Haraldsson, and John Woodward. Let's talk with developers, not about developers: A review of automatic program repair research. *IEEE TSE*, 2022.
- [22] Kai Huang, Zhengzi Xu, Su Yang, Hongyu Sun, Xuejun Li, Zheng Yan, and Yuqing Zhang. A survey on automated program repair techniques. *arXiv:2303.18184*, 2023.
- [23] Quanjun Zhang, Chunrong Fang, Yuxiang Ma, Weisong Sun, and Zhenyu Chen. A survey of learning-based automated program repair. *ACM TOSEM*, 2023.
- [24] Xin Zhou, Sicong Cao, Xiaobing Sun, and David Lo. Large language model for vulnerability detection and repair: Literature review and roadmap. *arXiv:2404.02525*, 2024.
- [25] Yunlong Xing, Shu Wang, Shiyu Sun, Xu He, Kun Sun, and Qi Li. What if is not enough? fixing null pointer dereference with contextual check. In *USENIX Security*, 2024.
- [26] Yinbo Yu, Yuanqi Xu, Kepu Huang, and Jiajia Liu. Tapfixer: Automatic detection and repair of home automation vulnerabilities based on negated-property reasoning. In *Proc. USENIX Security*, 2024.
- [27] Shiqi Shen, Aashish Kolluri, Zhen Dong, Prateek Saxena, and Abhik Roychoudhury. Localizing vulnerabilities statistically from one exploit. In *Proc. ACM ASIACCS*, 2021.

- [28] Yisroel Mirsky, George Macon, Michael Brown, Carter Yagemann, Matthew Pruett, Evan Downing, Sukarno Mertoguno, and Wenke Lee. Vulchecker: Graph-based vulnerability localization in source code. In *Proc. USENIX Security*, 2023.
- [29] Hiralal Agrawal and Joseph R Horgan. Dynamic program slicing. *ACM SIGPlan Notices*, 1990.
- [30] Prashant Hari Narayan Rajput, Constantine Doumanidis, and Michail Maniatakos. Iespach: Automated vulnerability localization and non-intrusive hotpatching in industrial control systems using data dependence graphs. In *Proc. USENIX Secur. Symp*, 2023.
- [31] Andreas Zeller and Ralf Hildebrandt. Simplifying and isolating failure-inducing input. *IEEE TSE*, 2002.
- [32] Dandan Xu, Di Tang, Yi Chen, XiaoFeng Wang, Kai Chen, Haixu Tang, and Longxing Li. Racing on the negative force: Efficient vulnerability root-cause analysis through reinforcement learning on counterexamples. In *USENIX Security*, 2024.
- [33] Jian Zhang, Chong Wang, Anran Li, Weisong Sun, Cen Zhang, Wei Ma, and Yang Liu. An empirical study of automated vulnerability localization with large language models. *arXiv:2404.00287*, 2024.
- [34] Shu Wang, Xinda Wang, Kun Sun, Sushil Jajodia, Haining Wang, and Qi Li. Graphspd: Graph-based security patch detection with enriched code semantics. In *Proc. IEEE S&P*, 2023.
- [35] Marcel Böhme, Van-Thuan Pham, Manh-Dung Nguyen, and Abhik Roychoudhury. Directed greybox fuzzing. In *Proc. ACM CCS*, 2017.
- [36] Zhiyuan Jiang, Shuitao Gan, Adrian Herrera, Flavio Toffalini, Lucio Romero, Chaojing Tang, Manuel Egele, Chao Zhang, and Mathias Payer. Evocatio: Conjuring bug capabilities from a single poc. In *Proc. ACM CCS*, 2022.
- [37] Lingchao Chen, Yicheng Ouyang, and Lingming Zhang. Fast and precise on-the-fly patch validation for all. In *IEEE/ACM ICSE*, 2021.
- [38] Luca Gazzola, Daniela Micucci, and Leonardo Mariani. Automatic software repair: A survey. *IEEE TSE*, 2017.
- [39] Qi Xin and Steven P Reiss. Leveraging syntax-related code for automated program repair. In *Proc. IEEE/ACM ASE*, 2017.
- [40] Xuan-Bach D Le, Duc-Hiep Chu, David Lo, Claire Le Goues, and Willem Visser. S3: syntax-and semantic-guided repair synthesis via programming by examples. In *Proc. ACM ESEC/FSE*, 2017.
- [41] Shin Hwei Tan, Hiroaki Yoshida, Mukul R Prasad, and Abhik Roychoudhury. Anti-patterns in search-based program repair. In *Proc. ACM/IEEE FSE*, 2016.
- [42] Wei Jin, Alessandro Orso, and Tao Xie. Automated behavioral regression testing. In *Proc. IEEE/ACM ICST*, 2010.
- [43] Qi Xin and Steven P Reiss. Identifying test-suite-overfitted patches through test case generation. In *Proc. ACM ISSTA*, 2017.
- [44] Kunal Taneja and Tao Xie. Diffgen: Automated regression unit-test generation. In *Proc. IEEE/ACM ASE*, 2008.
- [45] Dongsun Kim, Jaechang Nam, Jaewoo Song, and Sunghun Kim. Automatic patch generation learned from human-written patches. In *IEEE ICSME*, 2013.
- [46] Yingfei Xiong, Xinyuan Liu, Muhan Zeng, Lu Zhang, and Gang Huang. Identifying patch correctness in test-based program repair. In *Proc. IEEE/ACM ICSE*, 2018.
- [47] Hyungsub Kim, Muslum Ozgur Ozmen, Z Berkay Celik, Antonio Bianchi, and Dongyan Xu. Patchverif: Discovering faulty patches in robotic vehicles. In *Proc. USENIX Security*, 2023.
- [48] Jonathan Gallagher, Robin Gonzalez, and Michael E Locasto. Verifying security patches. In *Proceedings of the 2014 international workshop on privacy & security in programming*, 2014.
- [49] Xin Zhou, Bowen Xu, Kisub Kim, DongGyun Han, Hung Huu Nguyen, Thanh Le-Cong, Junda He, Bach Le, and David Lo. Leveraging large language model for automatic patch correctness assessment. *IEEE TSE*, 2024.
- [50] Nan Jiang, Thibaud Lutellier, and Lin Tan. Cure: Code-aware neural machine translation for automatic program repair. In *2021 IEEE/ACM ICSE*.
- [51] Xin Yin, Chao Ni, Shaohua Wang, Zhenhao Li, Limin Zeng, and Xiaohu Yang. Thinkrepair: Self-directed automated program repair. In *Proc. ACM ISSTA*, 2024.
- [52] Umair Z Ahmed, Pawan Kumar, Amey Karkare, Purushottam Kar, and Sumit Gulwani. Compilation error repair: for the student programs, from the student programs. In *Proc. ICSE*, 2018.
- [53] Ali Mesbah, Andrew Rice, Emily Johnston, Nick Glorioso, and Edward Aftandilian. Deepdelta: learning to repair compilation errors. In *Proc. ACM ESEC/FSE*, 2019.
- [54] Xueyang Li, Shangqing Liu, Ruitao Feng, Guozhu Meng, Xiaofei Xie, Kai Chen, and Yang Liu. Transrepair: Context-aware program repair for compilation errors. In *Proc. IEEE/ACM ASE*, 2022.
- [55] Westley Weimer. Patches as better bug reports. In *Proceedings of the 5th international conference on Generative programming and component engineering*, 2006.
- [56] Alex Shaw, Dusten Doggett, and Munawar Hafiz. Automatically fixing c buffer overflows using program transformations. In *Proc. IEEE DSN*, 2014.
- [57] Fengjuan Gao, Linzhang Wang, and Xuandong Li. Bovinspector: automatic inspection and repair of buffer overflow vulnerabilities. In *Proc. IEEE/ACM ASE*, 2016.
- [58] Zhen Huang, David Lie, Gang Tan, and Trent Jaeger. Using safety properties to generate vulnerability patches. In *Proc. IEEE S&P*, 2019.
- [59] Qing Gao, Yingfei Xiong, Yaqing Mi, Lu Zhang, Weikun Yang, Zhaoping Zhou, Bing Xie, and Hong Mei. Safe memory-leak fixing for c programs. In *IEEE/ACM ICSE*, 2015.
- [60] Seongjoon Hong, Junhee Lee, Jeongsoo Lee, and Hakjoo Oh. Saver: scalable, precise, and safe memory-error repair. In *Proc. ACM/IEEE ICSE*, 2020.
- [61] Hua Yan, Yulei Sui, Shiping Chen, and Jingling Xue. Automated memory leak fixing on value-flow slices for c programs. In *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, 2016.
- [62] Xi Cheng, Min Zhou, Xiaoyu Song, Ming Gu, and Jianguang Sun. Automatic integer error repair by proper-type inference. *IEEE TDSC*, 2019.
- [63] Paul Muntean, Martin Monperrus, Hao Sun, Jens Grossklags, and Claudia Eckert. Intrepair: Informed repairing of integer overflows. *IEEE TSE*, 2019.
- [64] Baojian Hua, Wanrong Ouyang, Chengman Jiang, Qiliang Fan, and Zhizhong Pan. Rupair: towards automatic buffer overflow detection and rectification for rust. In *Proc. ACM ASIACCS*, 2021.

- [65] Yuchi Tian and Baishakhi Ray. Automatically diagnosing and repairing error handling bugs in c. In *Proc. ACM/IEEE ICSE*, 2017.
- [66] Yuyao Zhang, Siqi Ma, Juanru Li, Kailai Li, Surya Nepal, and Dawu Gu. Smartshield: Automatic smart contract protection made easy. In *IEEE SANER*, 2020.
- [67] Michael Rodler, Wenting Li, Ghassan O Karame, and Lucas Davi. Evmpatch: Timely and automated patching of ethereum smart contracts. In *Proc. USENIX Security*, 2021.
- [68] Christof Ferreira Torres, Hugo Jonker, and Radu State. Elysium: Context-aware bytecode-level patching to automatically heal vulnerable smart contracts. In *Proceedings of the 25th International Symposium on Research in Attacks, Intrusions and Defenses*, 2022.
- [69] Tai D Nguyen, Long H Pham, and Jun Sun. Sguard: towards fixing vulnerable smart contracts automatically. In *Proc. IEEE S&P*, 2021.
- [70] Pengcheng Fang. Contractfix: A framework for automatically fixing vulnerabilities in smart contracts. *arXiv:2307.08912*, 2023.
- [71] Aharon Abadi, Ran Ettinger, Yishai A Feldman, and Mati Shomrat. Automatically fixing security vulnerabilities in java code. In *Proc. ACM OOPSLA*, 2011.
- [72] Xuezheng Xu, Yulei Sui, Hua Yan, and Jingling Xue. Vfix: value-flow-guided precise program repair for null pointer dereferences. In *Proc. ACM/IEEE ICSE*, 2019.
- [73] Yeting Li, Yecheng Sun, Zhiwu Xu, Jialun Cao, Yuekang Li, Rongchen Li, Haiming Chen, Shing-Chi Cheung, Yang Liu, and Yang Xiao. Regexscalpel: Regular expression denial of service (redos) defense by localize-and-fix. In *USENIX Security*, 2022.
- [74] Siqi Ma, David Lo, Teng Li, and Robert H Deng. Cdrep: Automatic repair of cryptographic misuses in android applications. In *Proc. ACM ASIACCS*, 2016.
- [75] Ying Zhang, Ya Xiao, Md Mahir Asef Kabir, Danfeng Yao, and Na Meng. Example-based vulnerability detection and repair in java code. In *Proc. IEEE/ACM ICPC*, 2022.
- [76] Siqi Ma, Ferdian Thung, David Lo, Cong Sun, and Robert H Deng. Vurle: Automatic vulnerability detection and repair by learning from examples. In *Proc. IEEE ESORICS*, 2017.
- [77] Gregory J Duck, Xiang Gao, and Abhik Roychoudhury. Binary rewriting without control flow recovery. In *Proc. ACM PLDI*, 2020.
- [78] Westley Weimer, ThanhVu Nguyen, Claire Le Goues, and Stephanie Forrest. Automatically finding patches using genetic programming. In *Proc. ACM/IEEE ICSE*, 2009.
- [79] Stephanie Forrest, ThanhVu Nguyen, Westley Weimer, and Claire Le Goues. A genetic programming approach to automated software repair. In *Proc. ACM/IEEE GECCO*, 2009.
- [80] Claire Le Goues, Michael Dewey-Vogt, Stephanie Forrest, and Westley Weimer. A systematic study of automated program repair: Fixing 55 out of 105 bugs for \$8 each. In *Proc. ACM/IEEE ICSE*, 2012.
- [81] Claire Le Goues, ThanhVu Nguyen, Stephanie Forrest, and Westley Weimer. Genprog: A generic method for automatic software repair. *IEEE TSE*, 2011.
- [82] Alexandru Marginean, Johannes Bader, Satish Chandra, Mark Harman, Yue Jia, Ke Mao, Alexander Mols, and Andrew Scott. Sapfix: Automated end-to-end repair at scale. In *Proc. ACM/IEEE ICSE*, 2019.
- [83] Ibrahim Mesecan, Daniel Blackwell, David Clark, Myra B Cohen, and Justyna Petke. Hypergi: Automated detection and repair of information flow leakage. In *Proc. IEEE/ACM ASE*, 2021.
- [84] Thomas Durieux, Benoit Cornu, Lionel Seinturier, and Martin Monperrus. Dynamic patch generation for null pointer exceptions using metaprogramming. In *Proc. IEEE SANER*, 2017.
- [85] Palina Tolmach, Yi Li, and Shang-Wei Lin. Property-based automated repair of defi protocols. In *Proc. IEEE/ACM ASE*, 2022.
- [86] Youkun Shi, Yuan Zhang, Tianhan Luo, Xiangyu Mao, Yinzhi Cao, Ziwen Wang, Yudi Zhao, Zongan Huang, and Min Yang. Backporting security patches of web applications: A prototype design and implementation on injection vulnerability patches. In *USENIX Security*, 2022.
- [87] Ridwan Shariffdeen, Xiang Gao, Gregory J Duck, Shin Hwei Tan, Julia Lawall, and Abhik Roychoudhury. Automated patch backporting in linux (experience paper). In *Proc. ACM ISSTA*, 2021.
- [88] Su Yang, Yang Xiao, Zhengzi Xu, Chengyi Sun, Chen Ji, and Yuqing Zhang. Enhancing oss patch backporting with semantics. In *Proc. ACM SIGSAC CCS*, 2023.
- [89] Ridwan Salihin Shariffdeen, Shin Hwei Tan, Mingyuan Gao, and Abhik Roychoudhury. Automated patch transplantation. *ACM TOSEM*, 2020.
- [90] Soeul Son, Kathryn S McKinley, and Vitaly Shmatikov. Fix me up: Repairing access-control bugs in web applications. In *NDSS*, 2013.
- [91] Zhengzi Xu, Yulong Zhang, Longri Zheng, Liangzhao Xia, Chenfu Bao, Zhi Wang, and Yang Liu. Automatic hot patch generation for android kernels. In *USENIX Security*, 2020.
- [92] Rijnard van Tonder and Claire Le Goues. Static automated program repair for heap properties. In *Proc. ACM/IEEE ICSE*, 2018.
- [93] Junhee Lee, Seongjoon Hong, and Hakjoo Oh. Memfix: static analysis-based repair of memory deallocation errors for c. In *Proc. ACM ESEC/FSE*, 2018.
- [94] Ruian Duan, Ashish Bijlani, Yang Ji, Omar Alrawi, Yiyuan Xiong, Moses Ike, Brendan Saltaformaggio, and Wenke Lee. Automating patching of vulnerable open-source software versions in application binaries. In *NDSS*, 2019.
- [95] Yu Liu, Sergey Mechtaev, Pavle Subotić, and Abhik Roychoudhury. Program repair guided by datalog-defined static analysis. In *Proc. ESEC/FSE*, 2023.
- [96] Nariyoshi Chida and Tachio Terauchi. Repairing dos vulnerability of real-world regexes. In *Proc. IEEE S&P*, 2022.
- [97] Hoang Duong Thien Nguyen, Dawei Qi, Abhik Roychoudhury, and Satish Chandra. Semfix: Program repair via semantic analysis. In *Proc. ACM/IEEE ICSE*, 2013.
- [98] Sergey Mechtaev, Jooyong Yi, and Abhik Roychoudhury. Angelix: Scalable multiline program patch synthesis via symbolic analysis. In *Proc. IEEE/ACM ICSE*, 2016.
- [99] Ridwan Shariffdeen, Yannic Noller, Lars Grunske, and Abhik Roychoudhury. Concolic program repair. In *Proc. ACM/IEEE ICSE*, 2021.
- [100] Yalin Ke, Kathryn T Stolee, Claire Le Goues, and Yuriy Brun. Repairing programs with semantic code search(t). In *Proc. IEEE/ACM ASE*, 2015.
- [101] Sergey Mechtaev, Jooyong Yi, and Abhik Roychoudhury. Directfix: Looking for simple program repairs. In *Proc. ACM/IEEE ICSE*, 2015.

- [102] Yue Chen, Yulong Zhang, Zhi Wang, Liangzhao Xia, Chenfu Bao, and Tao Wei. Adaptive android kernel live patching. In *Proc. USENIX Security*, 2017.
- [103] Xiang Gao, Sergey Mechtaev, and Abhik Roychoudhury. Crash-avoiding program repair. In *Proc. ACM ISSTA*, 2019.
- [104] Jifeng Xuan, Matias Martinez, Favio Demarco, Maxime Clement, Sebastian Lamelas Marcote, Thomas Durieux, Daniel Le Berre, and Martin Monperrus. Nopol: Automatic repair of conditional statement bugs in java programs. *IEEE TSE*, 2016.
- [105] Jacob Harer, Onur Ozdemir, Tomo Lazovich, Christopher Reale, Rebecca Russell, Louis Kim, et al. Learning to repair software vulnerabilities with generative adversarial networks. *Advances in neural information processing systems*, 2018.
- [106] Kai Huang, Su Yang, Hongyu Sun, Chengyi Sun, Xuejun Li, and Yuqing Zhang. Repairing security vulnerabilities using pre-trained programming language models. In *IEEE DSN-W*, 2022.
- [107] Nafis Tanveer Islam and Peyman Najafirad. Code security vulnerability repair using reinforcement learning with large language models. *arXiv:2401.07031*, 2024.
- [108] Fangzhou Wu, Qingzhao Zhang, Ati Priya Bajaj, Tiffany Bao, Ning Zhang, Ruoyu Wang, Chaowei Xiao, et al. Exploring the limits of chatgpt in software security applications. *arXiv:2312.05275*, 2023.
- [109] Baleegh Ahmad, Shailja Thakur, Benjamin Tan, Ramesh Karri, and Hammond Pearce. Fixing hardware security bugs with large language models. *arXiv:2302.01215*, 2023.
- [110] Ruoke Wang, Zongjie Li, Chaozheng Wang, Yang Xiao, and Cuiyun Gao. Navrepair: Node-type aware c/c++ code vulnerability repair. *arXiv:2405.04994*, 2024.
- [111] Peng Liu, He Wang, Chen Zheng, and Yuqing Zhang. Prompt fix: Vulnerability automatic repair technology based on prompt engineering. In *Proc. IEEE ICNC*, 2024.
- [112] Peiyu Liu, Junming Liu, Lirong Fu, Kangjie Lu, Yifan Xia, Xuhong Zhang, Wenzhi Chen, Haiqin Weng, Shouling Ji, and Wenhai Wang. Exploring chatgpt's capabilities on vulnerability management. In *USENIX Security*, 2024.
- [113] Michael Fu, Van Nguyen, Chakkrit Tantithamthavorn, Dinh Phung, and Trung Le. Vision transformer inspired automated vulnerability repair. *ACM TOSEM*, 2024.
- [114] Yu Nong, Mohammed Aldeen, Long Cheng, Hongxin Hu, Feng Chen, and Haipeng Cai. Chain-of-thought prompting of large language models for discovering and fixing software vulnerabilities. *arXiv:2402.17230*, 2024.
- [115] Arshiya Khan, Guannan Liu, and Xing Gao. Code vulnerability repair with large language model using context-aware prompt tuning. *arXiv:2409.18395*, 2024.
- [116] Che Wang, Jiashuo Zhang, Jianbo Gao, Libin Xia, Zhi Guan, and Zhong Chen. Contracttinker: Llm-empowered vulnerability repair for real-world smart contracts. In *Proc. IEEE/ACM ASE*, 2024.
- [117] Ummay Kulsum, Haotian Zhu, Bowen Xu, and Marcelo d'Amorim. A case study of llm for automated vulnerability repair: Assessing impact of reasoning and patch validation feedback. In *Proceedings of the 1st ACM International Conference on AI-Powered Software*, 2024.
- [118] Xin Zhou, Kisub Kim, Bowen Xu, DongGyun Han, and David Lo. Out of sight, out of mind: Better automatic vulnerability repair by broadening input ranges and sources. In *IEEE/ACM ICSE*, 2024.
- [119] Norbert Tihanyi, Ridhi Jain, Yiannis Charalambous, Mohamed Amine Ferrag, Youcheng Sun, and Lucas C Cordeiro. A new era in software security: Towards self-healing software via large language models and formal verification. *arXiv:2305.14752*, 2023.
- [120] Hugo Lefeuvre, Nathan Dautenhahn, David Chisnall, and Pierre Olivier. Sok: Software compartmentalization. In *Proc. IEEE S&P*, 2025.
- [121] Piergiorgio Ladisa, Henrik Plate, Matias Martinez, and Olivier Barais. Sok: Taxonomy of attacks on open-source software supply chains. In *Proc. IEEE S&P*, 2023.
- [122] Core rankings portal. <https://www.core.edu.au/conference-portal>, 2024. Access on: May, 2024.
- [123] Shreyas Sakharkar. Systematic review: Analysis of coding vulnerabilities across languages. *Journal of Information Security*, 2023.
- [124] Jennifer Gregory. Can memory-safe programming languages kill 70 <https://securityintelligence.com/news/memory-safe-programming-languages-security-bugs/>, 2024. Accessed on: May, 2024.
- [125] RIDWAN SALIHIN SHARIFFDEEN. Software vulnerability repair. 2022.
- [126] Jon Erickson. *Hacking: the art of exploitation*. 2008.
- [127] Cryptopedia Staff. What was the dao?, 2023.
- [128] Frances E Allen. Control flow analysis. *ACM Sigplan Notices*, 1970.
- [129] Jeanne Ferrante, Karl J Ottenstein, and Joe D Warren. The program dependence graph and its use in optimization. *ACM TOPLAS*, 1987.
- [130] Rust-ownerships and lifetimes. <https://doc.rust-lang.org/nomicon/ownership.html>. Accessed on June 2024.
- [131] Jayakrishna Vadayath, Moritz Eckert, Kyle Zeng, Nicolaas Weideman, Gokulkrishna Praveen Menon, Yanick Fratantonio, Davide Balzarotti, Adam Doupé, Tiffany Bao, Ruoyu Wang, et al. Arbiter: Bridging the static and dynamic divide in vulnerability discovery on binary programs. 2022.
- [132] Ralf Jung, Hoang-Hai Dang, Jeehoon Kang, and Derek Dreyer. Stacked borrows: an aliasing model for rust. *Proceedings of the ACM on Programming Languages*, (POPL), 2019.
- [133] Boqin Qin, Yilun Chen, Zeming Yu, Linhai Song, and Yiyang Zhang. Understanding memory and thread safety practices and issues in real-world rust programs. In *Proc. ACM PLDI*, 2020.
- [134] Wikipedia. Heuristic. [https://en.wikipedia.org/wiki/Heuristic_\(computer_science\)](https://en.wikipedia.org/wiki/Heuristic_(computer_science)), 2024. Accessed on: Jan, 2024.
- [135] Dawson Engler, David Yu Chen, Seth Hallem, Andy Chou, and Benjamin Chelf. Bugs as deviant behavior: A general approach to inferring errors in systems code. *ACM SIGOPS Operating Systems Review*, 2001.
- [136] Fuzzing. <https://en.wikipedia.org/wiki/Fuzzing>, 2024. Accessed on: Jan, 2024.
- [137] Shaobo He, Michael Emmi, and Gabriela Ciocarlie. ct-fuzz: Fuzzing for timing leaks. In *Proc. IEEE ICST*, 2020.
- [138] Metaprogramming. <https://en.wikipedia.org/wiki/Metaprogramming>, 2024. Accessed on: Jan, 2024.
- [139] Brenda S Baker. On finding duplication and near-duplication in large software systems. In *Proc. IEEE/ACM RE*, 1995.
- [140] Stéphane Ducasse, Matthias Rieger, and Serge Demeyer. A language independent approach for detecting duplicated code.

- In *Proc. IEEE ICSM*, 1999.
- [141] What is a patch? the essential guide for computer users. <https://www.lenovo.com/ie/en/glossary/what-is-alt-patch/>, 2024.
- [142] Nam H Pham, Tung Thanh Nguyen, Hoan Anh Nguyen, and Tien N Nguyen. Detection of recurring software vulnerabilities. In *Proc. IEEE/ACM ASE*, 2010.
- [143] Red Hat. Backporting security fixes. <https://access.redhat.com/security/updates/backporting>. Accessed on: Jan, 2024.
- [144] Yoann Padioleau, Julia Lawall, René Rydhof Hansen, and Gilles Muller. Documenting and automating collateral evolutions in linux device drivers. *ACM SIGOPS Operating Systems Review*, 2008.
- [145] Ferdian Thung, Xuan-Bach D Le, David Lo, and Julia Lawall. Recommending code changes for automatic backporting of linux device drivers. In *IEEE ICSE*, 2016.
- [146] Backports project. https://backports.wiki.kernel.org/index.php/Main_Page, 2024. Accessed on: Jan, 2024.
- [147] Xuan Bach D Le, David Lo, and Claire Le Goues. History driven program repair. In *2016 IEEE 23rd international conference on software analysis, evolution, and reengineering (SANER)*, 2016.
- [148] Ming Wen, Junjie Chen, Rongxin Wu, Dan Hao, and Shing-Chi Cheung. Context-aware patch generation for better automated program repair. In *IEEE ICSE*, 2018.
- [149] Tung Thanh Nguyen, Hoan Anh Nguyen, Nam H Pham, Jafar Al-Kofahi, and Tien N Nguyen. Recurring bug fixes in object-oriented programs. In *Proc. ACM/IEEE ICSE*, 2010.
- [150] Kim Marriott and Peter J Stuckey. *Programming with constraints: an introduction*. 1998.
- [151] Claire Le Goues, Michael Pradel, and Abhik Roychoudhury. Automated program repair. *Communications of the ACM*, 2019.
- [152] Stephen J Fink, Eran Yahav, Nurit Dor, G Ramalingam, and Emmanuel Geay. Effective tpestate verification in the presence of aliasing. *ACM TOSEM*, 2008.
- [153] Wikipedia. Exact cover. https://en.wikipedia.org/wiki/Exact_cover, 2024. Accessed on: Jan, 2024.
- [154] Datalog. <https://en.wikipedia.org/wiki/Datalog>, 2024. Accessed on: Jan, 2024.
- [155] James C King. Symbolic execution and program testing. *Communications of the ACM*, 1976.
- [156] Cristian Cadar, Daniel Dunbar, Dawson R Engler, et al. Klee: Unassisted and automatic generation of high-coverage tests for complex systems programs. In *Proc. OSDI*, 2008.
- [157] Patrice Godefroid, Nils Klarlund, and Koushik Sen. Dart: Directed automated random testing. In *Proceedings of the 2005 ACM SIGPLAN conference on Programming language design and implementation*, 2005.
- [158] Clark Barrett and Cesare Tinelli. *Satisfiability modulo theories*. 2018.
- [159] Konstantin Serebryany, Derek Bruening, Alexander Potapenko, and Dmitriy Vyukov. {AddressSanitizer}: A fast address sanity checker. In *2012 USENIX annual technical conference (USENIX ATC 12)*, 2012.
- [160] Clang. Undefinedbehaviorsanitizer. <https://clang.llvm.org/docs/UndefinedBehaviorSanitizer.html>, 2024. Accessed on: Jan, 2024.
- [161] Ricky W Butler. What is formal methods? *NASA LaRC Formal Methods Program*, 2001.
- [162] Ray J Solomonoff. A formal theory of inductive inference. part i. *Information and control*, 1964.
- [163] Wikipedia. Daikon (system). [https://en.wikipedia.org/wiki/Daikon_\(system\)](https://en.wikipedia.org/wiki/Daikon_(system)), 2024. Accessed on: May, 2024.
- [164] Haniel Barbosa, Clark Barrett, Martin Brain, Gereon Kremer, Hanna Lachnitt, Makai Mann, Abdalrhman Mohamed, Muthathir Mohamed, Aina Niemetz, Andres Nötzli, et al. cvc5: A versatile and industrial-strength smt solver. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, 2022.
- [165] Edmund M Clarke. Model checking. In *Foundations of Software Technology and Theoretical Computer Science: 17th Conference Kharagpur, India, December 18–20, 1997 Proceedings 17*, 1997.
- [166] Leonardo De Moura and Nikolaj Bjørner. Z3: An efficient smt solver. In *Proc. IEEE TACAS*, 2008.
- [167] cplusplus. strcpy function. <https://cplusplus.com/reference/cstring/strcpy/>, 2024. Accessed on: Jan, 2024.
- [168] Boolean satisfiability problem. https://en.wikipedia.org/wiki/Boolean_satisfiability_problem, 2024. Accessed on: Jan, 2024.
- [169] Divya Gopinath, Muhammad Zubair Malik, and Sarfraz Khurshid. Specification-based program repair using sat. In *Tools and Algorithms for the Construction and Analysis of Systems*, 2011.
- [170] Sergey Mechtaev, Alberto Griggio, Alessandro Cimatti, and Abhik Roychoudhury. Symbolic execution with existential second-order constraints. In *Proc. ACM ESEC/FSE*, 2018.
- [171] Common challenges faced in software requirements specification and how to overcome them. <https://www.practicallogix.com/common-challenges-faced-in-software-requirements-specification-and-how-to-overcome-them-your-ultimate-guide/>, 2024.
- [172] Fan Long and Martin Rinard. Automatic patch generation by learning correct code. In *Proc. ACM POPL*, 2016.
- [173] Ripon K Saha, Yingjun Lyu, Hiroaki Yoshida, and Mukul R Prasad. Elixir: Effective object-oriented program repair. In *Proc. IEEE/ACM ASE*, 2017.
- [174] Martin White, Michele Tufano, Matias Martinez, Martin Monperrus, and Denys Poshyvanyk. Sorting and transforming program repair ingredients via deep learning code similarities. In *Proc. IEEE SANER*, 2019.
- [175] Yi Li, Shaohua Wang, and Tien N Nguyen. Dlfix: Context-based code transformation learning for automated program repair. In *Proc. ACM/IEEE ICSE*, 2020.
- [176] Reudismam Rolim, Gustavo Soares, Loris D’Antoni, Oleksandr Polozov, Sumit Gulwani, Rohit Gheyi, Ryo Suzuki, and Björn Hartmann. Learning syntactic program transformations from examples. In *Proc. IEEE/ACM ICSE*, 2017.
- [177] Anil Koyuncu, Kui Liu, Tegawendé F Bissyandé, Dongsun Kim, Jacques Klein, Martin Monperrus, and Yves Le Traon. Fixminer: Mining relevant fix patterns for automated program repair. *Empirical Software Engineering*, 2020.
- [178] Rahul Gupta, Soham Pal, Aditya Kanade, and Shirish Shevade. Deepfix: Fixing common c language errors by deep learning. In *Proc. AAAI*, 2017.
- [179] Marko Vasic, Aditya Kanade, Petros Maniatis, David Bieber,

- and Rishabh Singh. Neural program repair by jointly learning to localize and repair. *arXiv:1904.01720*, 2019.
- [180] Zimin Chen, Steve Komrmusch, Michele Tufano, Louis-Noël Pouchet, Denys Poshyvanyk, and Martin Monperrus. Sequencer: Sequence-to-sequence learning for end-to-end program repair. *IEEE TSE*, 2019.
- [181] Michele Tufano, Cody Watson, Gabriele Bavota, Massimiliano Di Penta, Martin White, and Denys Poshyvanyk. An empirical study on learning bug-fixing patches in the wild via neural machine translation. *ACM TOSEM*, 2019.
- [182] Saikat Chakraborty, Yangruibo Ding, Miltiadis Allamanis, and Baishakhi Ray. Codit: Code editing with tree-based neural models. *IEEE TSE*, 2020.
- [183] Thibaud Lutellier, Hung Viet Pham, Lawrence Pang, Yitong Li, Moshi Wei, and Lin Tan. Coconut: combining context-aware neural translation models using ensemble for program repair. In *Proceedings of the 29th ACM SIGSOFT international symposium on software testing and analysis*, 2020.
- [184] Yi Li, Shaohua Wang, and Tien N Nguyen. Dear: A novel deep learning-based approach for automated program repair. *arXiv:2205.01859*, 2022.
- [185] Quanjun Zhang, Chunrong Fang, Bowen Yu, Weisong Sun, Tongke Zhang, and Zhenyu Chen. Pre-trained model-based automated software vulnerability repair: How far are we? *IEEE TDSC*, 2023.
- [186] Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 1996.
- [187] Shuo Ren, Daya Guo, Shuai Lu, Long Zhou, Shujie Liu, Duyu Tang, Neel Sundaresan, Ming Zhou, Ambrosio Blanco, and Shuai Ma. Codebleu: a method for automatic evaluation of code synthesis. *arXiv:2009.10297*, 2020.
- [188] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. *arXiv:1904.09675*, 2019.
- [189] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv:1707.06347*, 2017.
- [190] Erik Nijkamp, Hiroaki Hayashi, Caiming Xiong, Silvio Savarese, and Yingbo Zhou. Codegen2: Lessons for training llms on programming and natural languages. *arXiv:2305.02309*, 2023.
- [191] Codex. <https://openai.com/blog/openai-codex>, 2022.
- [192] Yue Wang, Weishi Wang, Shafiq Joty, and Steven CH Hoi. Codet5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation. *arXiv:2109.00859*, 2021.
- [193] Gpt-4. <https://openai.com/gpt-4>, 2023.
- [194] Mikhail R Gadelha, Felipe R Monteiro, Jeremy Morse, Lucas C Cordeiro, Bernd Fischer, and Denis A Nicole. Esmc 5.0: an industrial-strength c model checker. In *Proc. IEEE/ACM ASE*, 2018.
- [195] Paul E Black et al. Sard: A software assurance reference dataset. In *Anonymous Cybersecurity Innovation Forum.*, 2017.
- [196] Georgios Nikitopoulos, Konstantina Dritsa, Panos Louridas, and Dimitris Mitropoulos. Crossvul: a cross-language vulnerability dataset with commit data. In *Proc. ACM ESEC/FSE*, 2021.
- [197] Jingxuan He and Martin Vechev. Large language models for code: Security hardening and adversarial testing. In *Proc. ACM CCS*, 2023.
- [198] Guru Bhandari, Amara Naseer, and Leon Moonen. Cvefixes: automated collection of vulnerabilities and their fixes from open-source software. In *Proceedings of the 17th International Conference on Predictive Models and Data Analytics in Software Engineering*, 2021.
- [199] Serena Elisa Ponta, Henrik Plate, Antonino Sabetta, Michele Bezzi, and Cédric Dangremont. A manually-curated dataset of fixes to vulnerabilities of open-source software. In *Proc. IEEE/ACM MSR*, 2019.
- [200] Jiahao Fan, Yi Li, Shaohua Wang, and Tien N Nguyen. Ac/c++ code vulnerability dataset with code changes and cve summaries. In *Proc. IEEE/ACM MSR*, 2020.
- [201] Sofia Reis and Rui Abreu. A ground-truth dataset of real security patches. *arXiv:2110.09635*, 2021.
- [202] Quang-Cuong Bui, Riccardo Scandariato, and Nicolás E Díaz Ferreyra. Vul4j: a dataset of reproducible java vulnerabilities geared towards the study of program repair techniques. In *Proc. IEEE/ACM MSR*, 2022.
- [203] Paul E Black and Paul E Black. *Juliet 1.3 test suite*. 2018.
- [204] Eduard Pinconschi, Quang-Cuong Bui, Rui Abreu, Pedro Adão, and Riccardo Scandariato. Maestro: A platform for benchmarking automatic program repair tools on software vulnerabilities. In *Proc. ISSTA*, 2022.
- [205] Daniel Fried, Armen Aghajanyan, Jessy Lin, Sida Wang, Eric Wallace, Freda Shi, Ruiqi Zhong, Wen-tau Yih, Luke Zettlemoyer, and Mike Lewis. InCoder: A generative model for code infilling and synthesis. *arXiv:2204.05999*, 2022.
- [206] Jan Nowakowski Jan Keller. Ai-powered patching: the future of automated vulnerability fixes. Accessed on: May, 2024.
- [207] Nan Jiang, Kevin Liu, Thibaud Lutellier, and Lin Tan. Impact of code language models on automated program repair. In *IEEE/ACM ICSE*, 2023.
- [208] Google. Gemini. <https://gemini.google.com/>. Accessed on: May, 2024.
- [209] Itu-t recommendation t.81. <https://www.w3.org/Graphics/JPEG/itu-t81.pdf>, 1993. Accessed on: June, 2024.
- [210] Guangyuan Wu, Weining Cao, Yuan Yao, Hengfeng Wei, Taolue Chen, and Xiaoxing Ma. Llm meets bounded model checking: Neuro-symbolic loop invariant inference. In *Proc. IEEE/ACM ASE*, 2024.
- [211] Muhammad AA Pirzada, Giles Reger, Ahmed Bhayat, and Lucas C Cordeiro. Llm-generated invariants for bounded model checking without loop unrolling. In *Proc. IEEE/ACM ASE*, 2024.
- [212] Yuhang Wu, Zhenpeng Lin, Yueqi Chen, Dang K Le, Dongliang Mu, and Xinyu Xing. Mitigating security risks in linux with klaus: A method for evaluating patch correctness. In *Proc. USENIX Security*, 2023.
- [213] nfc: fix refcount leak in llcp_sock_connect(). <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=8a4cd82d62b5>, 2024.
- [214] Ruidi Qiu, Grace Li Zhang, Rolf Drechsler, Ulf Schlichtmann, and Bing Li. Correctbench: Automatic testbench generation with functional self-correction using llms for hdl design. *arXiv:2411.08510*, 2024.
- [215] Sofia Reis and Rui Abreu. Secbench: A database of real security vulnerabilities. In *SecSE@ ESORICS*, 2017.
- [216] Joern. <https://joern.io/>. Accessed on: May, 2024.

- [217] Github. Codeql. <https://codeql.github.com/>, 2024. Accessed on: Jan, 2024.
- [218] Justin Smith, Lisa Nguyen Quang Do, and Emerson Murphy-Hill. Why can't johnny fix vulnerabilities: A usability evaluation of static analysis tools for security. In *SOUPS*, 2020.
- [219] Omer Tripp, Salvatore Guarnieri, Marco Pistoia, and Aleksandr Aravkin. Aletheia: Improving the usability of static security analysis. In *ACM CCS*, 2014.

Appendix

The appendix illustrates more details on the experimental setting, dataset selection, mutation strategies, security patch generation artifacts status, and case studies.

A Dataset Selection Details

We selected the vulnerability and fixes dataset from the known datasets shown in Table 5. In these datasets, only SARD is synthetic data, only VJBench&VJbench-Trans, Vul4J, VulnLoc, and ExtractFix consist of test suites/exploits, so we select these datasets as our benchmarks. For \mathcal{D}_{SARD} , where defects and fixes are often associated with callee functions, caller functions' naming conventions, such as Goodxxx or Badxxx, may lead LLMs to erroneously transform "good" to "bad". To address this issue, we implemented an inter-procedural analysis with a depth of one, embedding callee functions' bodies and parameters within their caller functions (As shown in Figure 4).

Table 5: Vulnerability & Fixes Dataset Details

Name	Programming Language ¹	Data Type ²	CWE ³	Total ⁴	Patch ⁵	Test Suites/Exploit ⁶
SVEN[197]	Python&C	R	9	1,606	✓	✗
CVEfixes[198]	>30	R	209	5,365	✓	✗
SARD[195]	C/C++	S	60	25,297	✓	✗
Big-Vul[200]	C/C++	R	91	3,754	✓	✗
CrossVul[196]	>40	R	161	5,131	✓	✗
Vulas[199]	Java	R	NA	624	✗	✗
VJBench&VJbenchTrans[14]	Java	R	23	42	✓	✓
Vul4J[202]	Java	R	23	79	✓	✓
VulnLoc[27]	C/C++	R	12	43	✓	✓
ExtractFix[9]	C/C++	R	11	30	✓	✓
SecBench[215]	>20	R	51	676	✗	✗

¹ Programming Language: Indicates the types of programming languages covered by each dataset.

² Data Type: 'R' stands for real-world vulnerability and 'S' signifies synthetic data.

³ CWE: The number of CWE covered by the dataset. 'NA' signifies data not provided.

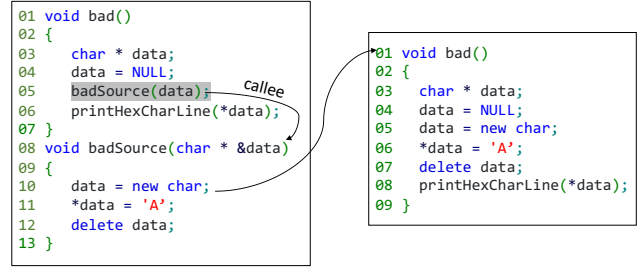
⁴ Total: The total number of entries or cases in the dataset.

⁵ Patch: Indicates whether patches are included in the dataset (✓/✗).

⁶ Test Suites/Exploit: Whether test suites or exploits are available (✓/✗).

B Experimental Setting Details

Experimental Environment Settings. Our experiments were conducted on Ubuntu 22.04.1 X86_64. To handle the varied dependencies needed for reproducing vulnerabilities, we employed Docker for consistent environments. We used GPT-4 and Gemini-Pro APIs and used CodeT5 and InCoder models and their fine-tuned models as Wu et al. [14]. We used Joern



(a) Example for SARD dataset with the root cause of vulnerability in callee function

(b) Refactored bad function

Figure 4: Inter-procedural processing for \mathcal{D}_{SARD}

(v2.0.110) [216] for \mathcal{D}_{SARD} preprocessing, CodeQL (v2.15.1) to detect remaining weakness.

Tools Applicability. Note that not all methods are applicable across benchmarks due to varying requirements (e.g., programming language) and target datasets. For instance, Vulnfix and ExtractFix require exploits, VRPilot needs exploit/test case feedback – these weren't used on the synthetic dataset. CodeT5 and its fine-tuned model were applied to Java only, as their training data lacked C/C++ datasets [207].

Evaluation Steps. We assessed the generated patches correctness through a three-step process. First, we checked compilation success. Second, we used an oracle mechanism: patches for vulnerabilities with existing tests (all real-world vulnerabilities) were verified against those tests, while synthetic data (\mathcal{D}_{SARD}) were checked using CodeQL [217]. Finally, all filtered patches were manually reviewed for semantic equivalence, functionality integrity, and absence of new vulnerabilities. To reduce subjectivity, two authors independently reviewed 200 randomly selected patches, achieving an agreement rate of 96.5% and a Cohen's kappa of 91.5%.

Input Configuration. For learning-based methods, we summarized the input configuration from [14, 15, 108, 114] and the models' usages. For deterministic models (CodeT5, InCoder, and their fine-tuned models), we only run them once; for non-deterministic models (i.e., the models that may produce different outputs for the same input, including GPT4 and Gemini here), we run them thrice, as long as one of the generated patches is correct, we regard them as correct. Especially, the practice prompts are derived from the above studies, though Nong et al. [114] highlighted the efficacy of CoT prompting. However, in the step of SPG, its core still lies in providing LLMs with the weakness information and the vulnerability location. So we keep providing CWE ID, CWE Name, and the buggy location for such generative models. For non-learning-based methods, we deploy the docker or virtual machine that the authors provided. The detailed input formats of learning-based methods are shown in Table 6.

C Mutation Strategies

To measure robustness on \mathcal{D}_{SARD} , we test three mutations:

- **Variable swapping:** Shuffle variable names within a func-

Table 6: Input configuration of learning-based methods

Model	Input
GPT-4-1106-preview	{"system": "You are a security patch generator. You will be given a vulnerable code, the CWE ID of it is {CWE_ID}, i.e.,{CWE_NAME}, the buggy line has been commented with <code>/*BUG*/</code> please directly provide the fixed code without any explanation.", "user": "vulnerable function"}
Gemini-pro	{"system_instruction": "You are a security patch generator. You will be given a vulnerable code, the CWE ID of it is {CWE_ID}, i.e.,{CWE_NAME}, the buggy line has been commented with <code>/*BUG*/</code> please directly provide the fixed code without any explanation.", "user": "vulnerable function"}
CodeT5-large	Mask buggy lines with <code><extra_id_i></code> and input the vulnerable function (i depends on the hunks(n) of modified code in the official patch, from 0 to n). In the input vulnerable function, also comment on buggy lines.
InCoder-6B	Mask buggy lines with and input the vulnerable function, (i depends on the hunks(n) of modified code in the official patch, from 0 to n-1). In the input vulnerable function, also comment on the buggy lines.
Fine-tuned CodeT5	The same with CodeT5.
Fine-tuned InCoder	The same with InCoder.

tion. For single variables, replace with a random string of equal length.

- **Condition structure reconstruction:** Switch if-else branches, e.g., change `if(a) {BlockA} else {BlockB}` to `if(!a) {BlockB} else {BlockA}`.
- **Loop transformation:** Transform loop constructs (e.g., for, while, do-while) to different structures.

D Additional Case Studies

```

1 @@ -287,7 +287,7 @@ fillpattern (int type, unsigned↵
    char *r, size_t size)
2     r[0] = (bits >> 4) & 255;
3     r[1] = (bits >> 8) & 255;
4     r[2] = bits & 255;
5 -   for (i = 3; i < size / 2; i += 2)
6 +   for (i = 3; i <= size / 2; i += 2)
7         memcpy (r + i, r, i);
8     if (i < size)
9         memcpy (r + i, r, size - i);

```

Listing 2: Patch for Gnubug-26545, repaired by learning-based methods but not by non-learning based methods

In Listing 2, the off-by-one error occurs due to using `<` instead of `<=` in the loop condition. For even size values, i could equal `size/2`, causing early loop exit. Changing to `<=` ensures the loop will execute when i equals `size/2`, fixing the vulnerability where the buffer wasn't fully filled under specific inputs. This vulnerability was not repaired by any non-learning-based methods because of missing constraints, whereas learning-based methods could infer from similar historical fixes.

Listing 3 is a single-hunk fix and only has two lines of code

addition. However, it cannot be repaired by both learning and non-learning methods. The root of the vulnerability lies in the reliance on the cached value of `td->td_nstrips`, which was originally computed to avoid redundant calculations. This value is calculated when the strip count is first needed, based on the image's length and the rows per strip. However, if the structure of the image changes afterward in another function, the cached value becomes outdated, leading to inconsistencies. The vulnerability emerges because there is no mechanism to ensure that the cached `td->td_nstrips` value is updated when the underlying image structure changes. For learning-based methods, it has no context, no information about the `td` member, and unclear about the root cause of this vulnerability, so it's hard to repair. For non-learning based methods, wrong constraints, and wrong localization lead to no successful repair. As analyzed in this case, although the change is 2 lines of code, it still has a complex logic.

```

1 @@ -5,6 +7,15 @@ TIFFNumberOfStrips(TIFF* tif)
2     TIFFDirectory *td = &tif->tif_dir;
3     uint32 nstrips;
4
5 +   if ( td->td_nstrips )
6 +       return td->td_nstrips;
7 +
8     nstrips = (td->td_rowsperstrip == (uint32) ↵
9                 -1 ? 1 :
10                TIFFHowmany_32(td->td_imagelength, td->↵
11                            td_rowsperstrip));
12     if (td->td_planarconfig == ↵
13         PLANARCONFIG_SEPARATE)

```

Listing 3: Patch for CVE-2016-9273, which was not repaired by both learning-based methods and non-learning based methods

E Future Directions

D8. Vulnerability Repair for Binaries. As we discussed in Section 3, most SPG focuses on source code-level fixes, which poses several limitations: (1) It requires access to the original build environment, which may be unavailable for older applications. (2) Developers may delay patching vulnerabilities in third-party libraries [77]. (3) Managing complex configurations and build options of open-source software is challenging [94]. To implement binary-level vulnerability fixes, we advocate integrate current repair methods with binary-rewriting or CodeLLMs for binary-level AVR.

D9. AVR Tools Usability. AVR can help reduce the workload of human. But how to better guide the use of the AVR tool effectively is still a challenging issue. Currently, most AVR tools are command-line based, which can present usability barriers for developers. Future research could also explore the AVR tools usability and develop usable AVR tools, previous work [218, 219] has stressed the usability of static security analysis tools but never for AVR tools.