# FINENIB: A QUERY SYNTHESIZER FOR STATIC ANALYSIS OF SECURITY VULNERABILITIES

#### **Anonymous authors**

000

001

002003004

010 011

012

013

014

016

018

019

021

024

025

026

027 028 029

031

033

034

036

040

041

042 043

044

046

047

048

049

051

052

Paper under double-blind review

#### **ABSTRACT**

Static analysis tools provide a powerful means to detect security vulnerabilities by specifying *queries* that encode vulnerable code patterns. However, writing such queries is challenging and requires diverse expertise in security and program analysis. To address this challenge, we present FineNib – an agentic framework that automatically synthesizes queries in CodeQL, a powerful static analysis engine, directly from a given CVE metadata. FineNib embeds an LLM in a synthesis loop with execution feedback, while constraining its reasoning using a custom MCP interface that allows structured interaction with a Language Server Protocol (for syntax guidance) and a RAG database (for semantic retrieval of queries and documentation). This approach allows FineNib to generate syntactically and semantically valid security queries. We evaluate FineNib on 176 existing CVEs across 111 Java projects. Building upon the Claude Code agent framework, FineNib synthesizes correct queries that detect the CVE in the vulnerable but not in the patched versions for 53.4% of CVEs. In comparison, using only Claude Code synthesizes 10% correct queries. Our generated queries achieve an F1 score of 0.7. In comparison, the general query suites in IRIS (a recent LLM-assisted static analyzer) and CodeQL only achieve F1 scores of 0.048 and 0.073, highlighting the benefit of FineNib's specialized synthesized queries.

#### 1 Introduction

Security vulnerabilities continue to grow at an unprecedented rate, with over 40,000 Common Vulnerabilities and Exposures (CVEs) reported in 2024 alone (CVE, 2025). Static analysis, a technique to analyze programs without executing them, is a common way of detecting vulnerabilities. Static analysis tools such as CodeQL (Github, 2025), Semgrep (Semgrep, 2023), and Infer (Meta, 2025) are widely used in industry. They provide domain-specific languages that allow specifying vulnerability patterns as queries. Such queries can be executed over structured representations of code, such as abstract syntax trees, to detect potential security vulnerabilities.

Despite their widespread use, existing query suites of static analysis tools are severely limited in coverage of vulnerabilities and precision. Extending them is difficult even for experts, as it requires knowledge of unfamiliar query languages, program analysis concepts, and security expertise. Incorrect queries can produce false alarms or miss bugs, limiting the effectiveness of static analysis. Correct queries can enable reliable detection of real vulnerabilities, supporting diverse use-cases such as regression testing, variant analysis, and patch validation, among others (Figure 1).

Meanwhile, CVE databases (MITRE, 2025; NIST, 2025; GitHub, 2025) provide rich information about security vulnerabilities, including natural language descriptions of vulnerability patterns and records of buggy and patched versions of the affected software repositories. This resource remains largely untapped in the automated construction of static analysis queries. Recent advances in LLMs, particularly in code understanding and generation, open up the possibility of leveraging this information to automatically synthesize queries from CVE descriptions, thereby bridging the gap between vulnerability reports and practical detection tools.

Synthesizing such queries poses significant challenges. The syntax of static analysis query languages is low-resource, richly expressive, and evolves continually. A typical query, such as the one in Figure 2(b) specifying a global dataflow pattern leaves ample room for errors in describing predicates for sources, sinks, sanitizers, and taint propagation steps. Even if the generated syntax is correct, success is measured by whether the query can identify at least one execution path travers-

056

060

061

062 063

064

065

066

067

068

069

071

073

074

075

076

077

079

081

083

084

085

087

090

092

094

096

098

099

100

101

102

103 104

105

106

107

Figure 1: A CodeQL query capturing a vulnerability pattern is synthesized by FineNib from an existing CVE and subsequently reused for regression testing, variant analysis, or patch validation.

ing the bug location in the vulnerable version while producing no matches in the patched version. Achieving this requires understanding the CVE context at the level of abstract syntax trees, such as code differences that introduce a sanitizer to prevent a flow from a source to a sink. Complicating matters further, reasoning about the code changes alone is often insufficient: sources, sinks, and taint propagation steps may reside in parts of the codebase far from the modified functions or files, and the vulnerability itself may involve non-trivial dataflow chains across these components. Thus, a correct query must not only integrate information from multiple locations across the program but also capture the intricate propagation patterns to accurately characterize the vulnerability.

In this paper, we present **FineNib** – an agentic framework that synthesizes queries in CodeQL, a powerful static analysis engine, directly from a given CVE metadata. We select CodeQL because it has the richest query language, which allows capturing complex inter-procedural vulnerability patterns. FineNib addresses the above challenges by embedding an LLM in a structured synthesis loop that incorporates execution feedback to verify query correctness and allows interactive reasoning using a custom MCP (Model Context Protocol) interface. The MCP interface constrains the model's reasoning using a Language Server Protocol (for syntax guidance) and a vector database of CodeQL queries and documentation (for semantic guidance). By combining these capabilities, FineNib avoids common pitfalls of naive LLM-based approaches, such as producing ill-formed queries, hallucinating deprecated constructs, or missing subtle vulnerability patterns, and instead produces queries that are both syntactically correct and semantically precise.

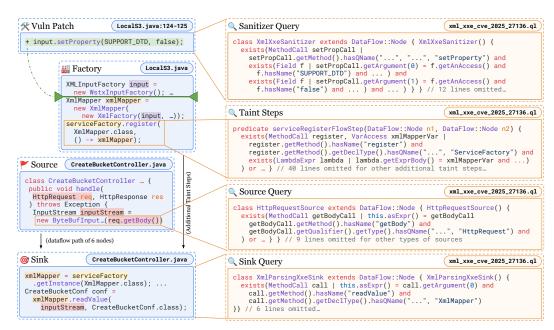
We evaluate FineNib on CWE-Bench-Java (Li et al., 2025b), which comprises 176 CVEs across 111 Java projects. These CVEs span 42 different Common Weakness Enumeration (CWE) categories and the projects range in size from 0.01 to 1.5 MLOC. To account for model training cut-offs, we include 65 CVEs reported during 2025 and target a recent CodeQL version 2.22.2 (July 2025). Using the Claude Code agent framework, FineNib achieves query compilation and success rates of 100% and 53.4%, compared to 19% and 0% for our best agentic baseline, Gemini CLI. Further, our generated queries have an F1 score of 0.7 for detecting true positive vulnerabilities, compared to 0.048 for IRIS (Li et al., 2025b), a recent LLM-assisted static analyzer, and 0.073 for CodeQL.

We summarize our main contributions:

- Agentic Framework for CVE-to-Query Synthesis. We present FineNib, an agentic framework that translates CVE descriptions into executable CodeQL queries, bridging the gap between vulnerability reports and static analysis. FineNib introduces a novel integration of execution-guided synthesis, semantic retrieval, and structured reasoning for vulnerability query generation.
- Evaluation on Real-World Repositories and CVEs. We evaluate FineNib on 176 CVEs in Java projects, covering 42 vulnerability types (CWEs) from CWE-Bench-Java. Each project involves complex inter-procedural vulnerabilities spanning multiple files. We show how FineNib can successfully identify sources, sinks, sanitizers, and taint propagation steps, and refine queries to ensure they raise alarms on vulnerable versions while remaining silent on patched versions.
- Comparison with Baselines. We compare FineNib against state-of-the-art agent frameworks and show that FineNib achieves substantially higher compilation, success, and F1 scores. We also compare FineNib's synthesized queries with state-of-the-art static analysis frameworks and show that our queries are more precise and have higher recall.

#### 2 ILLUSTRATIVE EXAMPLE

We illustrate the challenges of vulnerability query synthesis using CVE-2025-27136, an XML External Entity Injection (XXE) bug found in the repository Robothy/local-s3. Figure 2 depicts the vulnerability snippets, the patch, and the synthesized CodeQL query generated by FineNib.



- (a) The vulnerable dataflow snippets and the patch, which adds configuration to XMLInputFactory.
- (b) Snippets of the synthesized vulnerability query by FineNib capturing the patterns of dataflow source, sink, taint steps, and the sanitizer indicated by the vulnerability patch.

```
CodeQL Query

module HttpXxeFlow = TaintTracking::Global<HttpXxeFlowConfig>; // Config: HttpRequestSource, XmlParsingXxeSink, __
import HttpXxeFlow::PathNode source, HttpXxeFlow::PathNode sink
from HttpXxeFlow::PathNode source, HttpXxeFlow::PathNode sink
where HttpXxeFlow::flowPath(source, sink)
select sink.getNode(), source, sink, "HTTP request data flows to XML parser...", source.getNode(), "HTTP request entry point"
```

(c) The synthesized CodeQL path query that ties everything together.

Figure 2: Illustration of vulnerability CVE-2025-27136 in repository Robothy/local-s3 which exhibits an XML External Entity Injection weakness (CWE-611). When the XmlMapper is not configured to disable *Document Type Definition* (DTD), the function readValue may declare additional entities, allowing hackers to inject malicious behavior.

**Vulnerability context.** The vulnerability arises when the XmlMapper object is used to parse user-provided XML data (Figure 2a). In the vulnerable code, XmlMapper.readValue is called on the HTTP request body without disabling support for *Document Type Definitions* (DTDs). As a result, an attacker can inject malicious external entity declarations into the input stream, enabling server-side request forgery (SSRF) attacks, allowing for access to resources that should not be accessible from external networks, effectively leaking sensitive information. The patch mitigates the issue by configuring the underlying XMLInputFactory with the property SUPPORT\_DTD=false.

Synthesizing the query. The CodeQL query that can effectively capture the vulnerability pattern needs to incorporate 1) sources such as <code>HttpRequest.getBody</code> calls where untrusted malicious information enters the program, 2) sinks such as invocations of <code>XmlMapper.readValue</code>, where the XXE vulnerability is manifested, 3) additional taint steps related to how the <code>XmlMapper</code> is constructed and configured, involving non-trivial interprocedural flows spanning multiple files, and 4) sanitizers such as calls to <code>setProperty(SUPPORT\_DTD, false)</code>, so that we know that no alarm should be reported after the vulnerability has been fixed.

In general, the synthesized query must connect all these components to be able to detect the bug in the vulnerable program, while not reporting the same alarm after the vulnerability has been fixed. Figure 2b shows all the components of the CodeQL query (simplified), capturing their individual syntactic patterns. Lastly, Figure 2c connects all these components into a coherent path query by using CodeQL's TaintTracking::Global<.>::PathGraph and the SQL-like from-where-select query, which returns the exact path from source to sink.

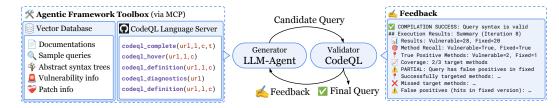


Figure 3: Overall pipeline of FineNib's iterative synthesis loop between an agentic query generator and a CodeQL-based validator. The generator uses a vector database and our CodeQL Language Server as tools while the validator produces compilation, execution, and coverage feedback.

**Challenges and solutions.** Vulnerability query synthesis must overcome several tightly-coupled challenges. We hereby state the challenges and explain how FineNib addresses them.

- Rich expressiveness and fragility of syntactic patterns. CodeQL is powerful but syntactically intricate: small mistakes in predicate names, qualifiers, or AST navigation often produce syntactically valid yet semantically useless queries. FineNib mitigates this fragility through its Language Server Protocol (LSP) interface for syntax guidance and RAG database for semantic retrieval of existing CodeQL queries and documentation. These structured interactions guide predicate selection and AST navigation during synthesis, reducing off-by-name and version-mismatch errors.
- Inter-procedural taint propagation across a large codebase. Sources, sinks, and sanitizers typically live in different modules or files and are connected by nontrivial inter-procedural flows (lambdas, factory patterns, etc.). While CodeQL provides robust inter-procedural analysis for many common patterns, gaps in dataflow still require bridging via additional taint propagation steps. Through its custom MCP interface, FineNib performs structured reasoning to discover candidate program points, synthesize custom taint-step predicates (e.g., service registration), and compose them into a CodeQL path query that tracks data across file and component boundaries.
- Semantic precision: alarm on the vulnerable version, silence on the patched version. A useful vulnerability query must not only parse correctly but also be discriminative. FineNib enforces this semantic requirement directly during synthesis. Via an iterative refinement loop, the successful criteria states that in the fixed program, there should be no alarm being raised about the vulnerability. This incentivizes the agent to synthesize sanitizer predicates (e.g., the setProperty call) and use them to constrain the path query so that sanitizer presence suppresses the alarm. The resulting query thus captures the exact behavior difference, producing alarms on the vulnerable snapshot and not on the patched snapshot.

Together, these capabilities let FineNib synthesize a semantically precise CodeQL query that can be reused for regression testing, variant analysis, or patch validation. We now elaborate on the detailed design and implementation of FineNib.

#### 3 FINENIB

At a high level, FineNib operates inside a repository-aware iterative refinement loop (Figure 3). In each iteration, the agent proposes a candidate CodeQL query, a CodeQL-based validator executes and scores it on both the vulnerable and patched versions of the repository, and the agent uses the validation feedback to propose targeted repairs. The loop terminates successfully when the validator accepts a query, or fails after a fixed iteration budget. In this section, we elaborate the major design components that make the loop effective.

#### 3.1 PROBLEM STATEMENT

The task of vulnerability detection is generally framed as a taint analysis task, where the goal of a *query* is to find dataflow paths from a *source* (e.g., an API endpoint accepting user input) to a *sink* (e.g., a database write) that lack proper *sanitization* (e.g., filtering malicious data).

We formalize the *Vulnerability Query Synthesis* problem as follows. Assume as input a vulnerable project version  $P_{\text{vuln}}$ , its fixed version  $P_{\text{fixed}}$ , and a textual CVE description (commonly available in open vulnerability reports). Let us assume we have inter-procedural dataflow program graphs for each code version:  $G_{\text{vuln}} = (V_{\text{vuln}}, E_{\text{vuln}})$  and  $G_{\text{fixed}} = (V_{\text{fixed}}, E_{\text{fixed}})$ . Let  $\Delta P$  denote the source-level patch between  $P_{\text{vuln}}$  and  $P_{\text{fixed}}$ . We represent the patch in the dataflow-graph domain

as a patch subgraph  $\Delta G=(\Delta V,\Delta E)$ , where  $\Delta V$  is the set of graph nodes that correspond to the modified program snippets.

A vulnerability path query Q evaluated on a graph G returns a set of dataflow paths, denoted as  $\Pi = \llbracket Q \rrbracket(G)$ . We write each path  $\pi \in \Pi$  as  $\pi = \langle v_1, \ldots, v_k \rangle$ , where each  $v_i \in V$  is a node in the dataflow graph G. Consecutive nodes  $(v_i, v_{i+1})$  should be either connected by an existing edge in E, or an additional taint step specified in the query Q, to compensate for missing edges via dataflow graph construction. Specifically, we call  $v_1$  the source of path  $\pi$  and  $v_k$  the sink of  $\pi$ .

Synthesis task. We aim to synthesize a query Q from the vulnerability report satisfying the following requirements:

- 1. Well-formedness. Q is syntactically valid (based on the latest CodeQL syntax) and can be executed on the target CodeQL infrastructure (e.g., dataflow graphs) without runtime errors.
- 2. Vulnerability detection. Q generates at least one path  $\pi$  in the vulnerable version that traverses the patched region:

$$\exists \pi \in \llbracket Q \rrbracket (G_{\mathrm{vuln}}) \quad \text{such that} \quad \pi \cap \Delta V \neq \emptyset.$$

3. **Fix discrimination.** *Q* does not report the vulnerability in the fixed version. Concretely, no path reported on the fixed version should traverse the patched locations:

$$\forall \pi \in [\![Q]\!](G_{\text{fixed}}), \text{ we have } \pi \cap \Delta V = \emptyset.$$

In other words, the synthesized query must be executable, must witness the vulnerability in the vulnerable version via a path that uses code touched by the fix, and must not attribute the same (patched) behavior in the fixed version. When only the well-formedness condition is satisfied, we say that the query Q is valid (denoted as  $\mathtt{valid}(Q)$ ); when all the conditions are satisfied, the query Q is successful (denoted as  $\mathtt{success}(Q; P_{\mathtt{vuln}}, P_{\mathrm{fixed}})$ ). Note that these criteria may admit potentially false positive paths in both versions. It might be possible to consider additional constraints regarding precision, but it might further complicate synthesis. In practice, we find most queries synthesized by FineNib already have high precision.

#### 3.2 DESIGN OF FINENIB

Concretely, FineNib proceeds in an iterative refinement loop indexed by  $i = [0, 1, \dots]$ . Via prompting, the LLM agent-based synthesizer first proposes an initial candidate query  $Q_0$ . For each iteration i, the validator evaluates  $Q_i$  and produces a feedback report. We consider synthesis successful at iteration i iff  $\mathtt{success}(Q_i; P_{\mathrm{vuln}}, P_{\mathrm{fixed}})$  holds; in that case the loop terminates and  $Q_i$  is returned. Otherwise, the synthesizer analyzes the feedback and the previous candidate  $Q_i$ , and produces the next query candidate  $Q_{i+1}$ . The loop stops  $\mathtt{successfully}$  when  $\mathtt{success}(\cdot)$  is achieved or fails once i reaches the pre-configured limit N (in our implementation N=10). The remainder of the design focuses on two aspects: 1) how the agentic synthesizer performs synthesis, and 2) how the validator generates and communicates feedback. We elaborate on both below.

**Agentic synthesizer.** In each iteration i, the LLM-based agentic synthesizer runs an inner *conversation loop* of up to M turns. In each turn, the agent either performs internal reasoning or issues a tool call by emitting a JSON-formatted action. When a tool call succeeds, the tool returns a JSON-formatted response that is appended to the conversation history. Conversation histories are kept local to the current refinement iteration (i.e., not carried over between iterations) to keep context compact and relevant. In practice, we set M=50, i.e., the agent may interact with tools up to 50 times before generating a candidate query for validation.

Two design choices are critical for the effectiveness of this loop: 1) the *initial prompt* that initializes and constrains the agent's behavior, and 2) the *toolbox* of callable tools, each exposed by a custom Model Context Protocol (MCP) server. We refer to the combined problem of designing these items as *Context Engineering* (discussed in Section 3.3).

**CodeQL-Based Validator.** The validator compiles and executes each candidate query against the vulnerable and fixed versions and returns a concise, structured feedback report that is used to drive refinement (Figure 3). The report contains: (i) CodeQL compilation results, (ii) execution counts (matches on vulnerable and fixed graphs), (iii) recall and coverage statistics, (iv) concrete counterexample traces and hit locations, and (v) a prioritized set of next-step recommendations (e.g., add qualifiers, synthesize sanitizer checks, or expand taint steps) that are programmatically generated via a template.

271

272

273

274

275

276

277

278

279

281

282

283 284

285

286

287

288

289

290

291

293

295

296

297298

299

300

301

302

303

304 305

306

307

308

309

310 311

312

313

314

315

316

317

318

319

320 321

322

323

Figure 4: Illustration of example traces of conversation during the synthesis of the query in the motivating example (Figure 2). LLM-agent may think, invoke tools that are available in the toolbox, and receive responses from the MCP servers.

#### 3.3 CONTEXT ENGINEERING FOR AGENTIC SYNTHESIZER

The primary goal of context engineering is to expose the LLM-based agent to the most *precise* amount of information: enough for the agent to make progress, but not so much that the LLM is confused or the cost explodes. As illustrated in Figure 3, FineNib relies on two primary MCP servers to provide demand-driven, structured information to the agent: a retrieval-augmented vector database and a CodeQL Language Server interface. We show example traces of conversation loop in Figure 4 and describe the available tools below.

**Initial prompt.** Each refinement iteration begins with an *initial prompt* that kicks-off the agentic conversation loop. The initial prompt in the first iteration contains a query skeleton for reference (See § A.1 for an example). In subsequent iterations, the prompt contains a summary of the synthesis goal and constraints, the previous candidate query  $Q_{i-1}$ , and the validator feedback report. Concretely, the initial prompt emphasizes: (i) the success predicate (see Success(·)), (ii) concrete counterexamples from previous feedback, and (iii) an explicit list of callable tools and their purpose.

**Vector database.** We use a retrieval-augmented vector database (ChromaDB MCP server in our implementation) to store large reference corpora without polluting the LLM prompt. The database is pre-populated with (i) vulnerability analysis notes and diffs, (ii) Common Weaknesses Enumeration (CWE) definitions, (iii) same-version CodeQL API documentation, (iv) curated CodeQL sample queries, and (v) small abstract syntax tree (AST) snippets extracted from the target repository. During a conversation loop, the agent issues compact retrieval queries (e.g., to fetch example CodeQL queries related to the CWE) and receives ranked documents or snippets on demand.

In practice, we may populate our RAG database with tens of thousands of documents. Even with this large corpus, we observe that the LLM-agent reliably retrieves exactly the kinds of artifacts it needs: CodeQL sample queries that inspire overall query structure, small AST snippets that suggest the precise syntactic navigation, and vulnerability writeups or diff excerpts that help discriminate buggy from patched behavior. These demand-driven lookups let the agent gather high-quality information without loading the main prompt with large reference corpora.

CodeQL language server. We expose the CodeQL Language Server (Github, d) through a MCP server that the agent can call for precise syntax-aware guidance. Importantly, we developed our own CodeQL Language Server client and MCP server that ensures syntactic validity (especially for the given CodeQL version) during query generation. The LLM agent's MCP client makes the tool call which is received by the CodeQL MCP server. The MCP server forwards tool calls, such as complete(file, loc, char), diagnostics(file), and definition(file, loc, char), to the underlying CodeQL process and returns JSON-serializable responses. Tools such as completion help the agent fill query templates and discover correct API or AST names, while diagnostics reveal compile or linter errors (e.g., unknown predicate names) that guide mutation. Appendix B shows the full specification and example request and response schemas.

#### 3.4 DISCUSSION: ALTERNATIVE DESIGNS

We discuss several alternative designs that we considered but found ineffective in practice. Allowing the agent unrestricted access to compile-and-run CodeQL via MCP led to severe performance degra-

Table 1: FineNib Query Success by CWE Type

CWE Type	<b>Total CVEs</b>	# Success	Success (%)	Avg Precision
CWE-022 (Path Traversal)	48	31	64.6	0.75
CWE-079 (Cross-Site Scripting)	36	18	50.0	0.621
CWE-094 (Code Injection)	20	12	60.0	0.606
CWE-078 (OS Command Injection)	12	7	58.3	0.628
CWE-502 (Deserialization)	6	4	66.7	0.853
CWE-611 (XXE)	5	3	60.0	0.657
Other CWEs ( $\leq$ 4 CVEs)	49	19	38.8	0.504
Total	176	94	53.4	0.631

dation: compilation and full execution are expensive operations that the LLM soon overused, so we instead expose only lightweight diagnostics during the conversation and defer full compile-and-run to the end of each iteration. Permitting free online search for vulnerability patterns or snippets similarly proved problematic. It is both costly and easy for the agent to rely on web lookups, which quickly pollutes the working context and degrades synthesis quality. Equipping the agent with an extensive set of heterogeneous tools led to confusion and poor tool-selection behavior; in contrast, a small, well-scoped toolbox yields more reliable actions. Finally, retaining full conversation histories across refinement iterations induced context rot and ballooning prompt sizes, so we keep histories local to each iteration. Overall, our current design is a pragmatic trade-off that balances cost, responsiveness, and synthesis effectiveness.

## 

#### 4 EVALUATION

We answer the following research questions in this work:

- **RQ 1**: For how many CVEs, can FineNib generate queries successfully?
- **RQ 2**: How useful is each component of FineNib?
- RQ 3: How does the choice of base agent framework affect FineNib's effectiveness?

#### 4.1 EXPERIMENTAL SETUP

We develop FineNib on top of the Claude Code framework (Anthropic, 2025) and use Claude Sonnet 4 for all our experiments. For agent baselines, we select Codex with GPT-5 (minimal reasoning) and Gemini CLI with Gemini 2.5 Flash. For each CVE and agent baseline, we use a maximum of 10 iterations (N=10). For static analysis baselines, we select IRIS Li et al. (2025b) and CodeQL (version 2.22.2) query suites. Experiments were run on machines with the following specifications: an Intel Xeon Gold 6248 2.50GHz CPU, four GeForce RTX 2080 Ti GPUs, and 750GB RAM.

**Dataset.** We used CWE-Bench-Java (Li et al., 2025b) and its latest update, which added new CVEs from 2025. We were able to successfully build and use 111 (out of 120) Java CVEs evaluated in IRIS (Li et al., 2025b), and 65 (out of 91) 2025 CVEs. Each sample in CWE-Bench-Java comes with the CVE metadata and fix commit information associated with the bug.

#### 4.2 EVALUATION METRICS

Besides Valid(Q) and  $Success(Q; P_{vuln}, P_{fixed})$ , we use the following terms and metrics when evaluating FineNib and baselines on the problem of vulnerability query synthesis:

$$\begin{split} \operatorname{Rec}(Q) &= \mathbbm{1} [\exists \pi \in [\![Q]\!] (G_{\operatorname{vuln}}), \pi \cap \Delta V \neq \emptyset], \quad \operatorname{Prec}(Q) = \frac{|\{\pi \in [\![Q]\!] (G_{\operatorname{vuln}}) \mid \pi \cap \Delta V \neq \emptyset\}|}{|[\![Q]\!] (G_{\operatorname{vuln}})|}, \\ & \operatorname{F1}(Q) = 2 \cdot \frac{\operatorname{Prec}(Q) \cdot \operatorname{Rec}(Q)}{\operatorname{Prec}(Q) + \operatorname{Rec}(Q)}. \end{split}$$

#### 4.3 RQ1: FINENIB EFFECTIVENESS

**FineNib vs. state-of-the-art QL.** Table 1 shows FineNib's overall query synthesis success rate by CWE. Table 2 shows the notable increase in precision between CodeQL, IRIS, and FineNib. We have successfully synthesized 53.4% of the CVEs. For half the queries FineNib correctly synthesized CodeQL, detect the CVE, did not find false positives when executed on the fixed version of the CVE's repository. The lack of true positive recall is why CodeQL and IRIS have significantly

Table 2: Recall Performance Comparison Across Methods (Shared CVEs: 130)

380
381
382
000



<u> </u>				
Method	Recall Rate (%)	Avg Precision	Avg F1 Score	
CodeQL	20.0	0.055	0.073	
IRIS	35.4	0.031	0.048	
FineNib	80.0	0.672	0.700	

Recall Rate Comparison by CWE Type Across Different Methods (Shared CVEs: 102)

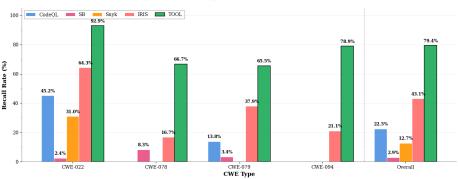


Figure 5: Recall Rate Comparison by CWE Type Across Different Methods (102 CVEs).

lower precision. CodeQL's queries are broad, categorized by CWE queries. IRIS generates all of the predicates for potential sources and sinks wit CodeQL, and does not generate sanitizer or taint step predicates.

**Impact of training cut-off.** We also want to take note that Claude Sonnet 4's training cut-off is March 2025. Table 3 shows that FineNib performs consistently regardless of CVEs before or after the cut-off period. The CodeQL version, 2.22.2, was released in July 2025. New versions of CodeQL often include analysis improvements and new QL packs (Github, a).

Table 3: Tool Performance Before vs After Training Cutoff

CVE Period	<b>Total CVEs</b>	# Recall	Success (%)	Avg Precision	Avg F1 Score
Pre-2025 (2011-2024)	111	64	57.7	0.676	0.702
2025+ (Post-cutoff)	65	30	46.2	0.555	0.583
Overall	176	94	53.4	0.631	0.658

#### 4.4 RQ2: ABLATION STUDIES

For ablations, we chose 20 CVEs and ran FineNib with one of the FineNib components removed (Table 4). The ablation with no tools refers to only running Claude Code with the iterative feedback system. The high recall rate when removing access to the AST cache while lowered recall rates without the LSP server or documentation access show that the LSP and documentation lookup impact the synthesis performance more. We also include FineNib's performance on the same set of CVEs, and point out its significantly higher query success rate and precision score. Claude Code without tools scored a high recall rate, yet failed to synthesize queries without false positives when executed on the fixed version.

#### 4.5 RO3: STATE OF THE ART AGENT COMPARISON

FineNib can be transferred to other coding agents by MCP configuration changes, since its tools are MCP servers. Changing agents involves using a different CLI command to start the coding agent, which was a minor adjustment for using FineNib. We used Gemini CLI with Gemini 2.5 Flash, and Codex with GPT-5 minimal. We evaluated their performance on 20 CVEs, and although there were no successful queries, we achieved an increase in compilation success for both agents compared to using the agents without FineNib in Table 5.

432

433 434

438 439 440

445 446

448 449 450

452 453

454

455

461

462

467 468 469

471 472 473

470

474 475

476 477 478

479 480

481

482 483 484

485

Table 4: Ablation Study (out of 20 CVEs)

Variant	% Successful	Recall Rate	Avg Precision	Avg F1 Score
FineNib	55%	80%	0.67	0.69
w/o LSP	25% (-30%)	55% (-25%)	0.32	0.36
w/o Doc/Ref	20% (-35%)	55% (-25%)	0.32	0.36
w/o AST	25% (-30%)	$80\% \ (\pm 0\%)$	0.41	0.47
w/o Tools	10% (-45%)	55% (-25%)	0.33	0.36

Table 5: LLM-agent baselines' compilation rate on 20 CVEs.

Agent Baselines	Configuration	Compilation Rate (%)
Gemini CLI	Gemini 2.5 Flash (with tools)	24
Gemini CLI	Gemini 2.5 Flash (without tools)	19
Codex	GPT 5 (with tools)	24
Codex	GPT 5 (without tools)	0

#### RELATED WORK

**LLMs** and vulnerability detection. LLMs have been used extensively for vulnerability detection and repair using techniques such as fine-tuning and prompt engineering (Zhou et al., 2024). LLMs have also been combined with existing program analysis tools for vulnerability detection. The combination of LLMs can be used from vulnerability analysis like IRIS's (Li et al., 2025b) source and sink identification, however IRIS depends on a limited set of CWE templates derived from CodeQL's CWE queries. IRIS also only the LLM for identifying sources and sinks. KNighter synthesizes CSA checkers given a fix commit of a C repository (Yang et al., 2025), however the checkers are written in C which has more available training data. MocQ's uses an LLM to derive a subset DSL of CodeQL and Joern, and then provides a feedback loop to the LLM though prompting via API calls is used rather than an agent with tools and MocQ uses significantly higher iterations, with a max threshold of 1,000 iterations per vulnerability experiment. (Li et al., 2025a).

**LLM agents and tool usage.** SWE-agent pioneered the idea of autonomous LLM agents using tools for software engineering tasks Yang et al. (2024). LSPAI Go et al. (2025), an IDE plugin, uses LSP servers to guide LLM-generated unit tests. Hazel, a live program sketching environment, uses a language server (Blinn et al., 2024) to assist code completions synthesized by LLMs. The Hazel Language Server provides the typing context of a program hole to be filled.

Low resource LLM code generation. SPEAC uses ASTs combined with constraint solving to repair LLM-generated code for low resource programming languages (Mora et al., 2024). SPEAC converts a buggy program into an AST and uses a solver to find the minimum set of AST nodes to replace, to satisfy language constraints. MultiPL-T generates datasets for low resource languages by translating high resource language code to the target language and validates translations with LLM generated unit tests (Cassano et al., 2024).

#### CONCLUSION AND LIMITATIONS

We present FineNib, an agentic framework for synthesizing syntactically correct and precise CodeQL queries given known vulnerability patterns. We will also open source our CodeQL LSP MCP server and FineNib. In future work, we plan to explore efficient ways to synthesize, and to combine our synthesized queries with dynamic analysis tools.

**Limitations.** We omit CVEs where the vulnerability involves non-Java code such as configuration files or other languages. FineNib can be used with exploit generation to find vulnerabilities that are realized during dynamic execution. For supporting other languages that can be queried by CodeQL, the vector database can be filled with references, documentation, and example queries in other CodeQL supported languages. We also want to note that Claude Sonnet 4's official training cut-off is March 2025, however the 2025 CVEs evaluated were reported between January to August 2025.

### REFERENCES

486

487 488

489

495

496

497

498 499

500

501

502

504

505

506

507

508 509

510

511

512

513 514

515

522

523

524

525

527

529

535

538

- Anthropic. Claude code: Agentic coding assistant. https://github.com/anthropics/claude-code, 2025. Version 1.0. Released February 2025. Accessed September 2025.
- Andrew Blinn, Xiang Li, June Hyung Kim, and Cyrus Omar. Statically Contextualizing Large Language Models with Typed Holes. Artifact for Statically Contextualizing Large Language Models with Typed Holes, 8(OOPSLA2):288:468–288:498, October 2024. doi: 10.1145/3689728. URL https://dl.acm.org/doi/10.1145/3689728.
  - Federico Cassano, John Gouwar, Francesca Lucchetti, Claire Schlesinger, Anders Freeman, Carolyn Jane Anderson, Molly Q Feldman, Michael Greenberg, Abhinav Jangda, and Arjun Guha. Knowledge transfer from high-resource to low-resource programming languages for code llms, 2024. URL https://arxiv.org/abs/2308.09895.
  - CVE. Common vulnerabilities and exposures (cve). https://www.cve.org/about/Metrics, 2025.
  - Github. CodeQL 2.23.0 (2025-09-04) CodeQL, a. URL https://codeql.github.com/docs/codeql-overview/codeql-changelog/codeql-cli-2.23.0/.
  - Github. codeql/java/ql/lib/semmle/code/java/security at main · github/codeql, b. URL https://github.com/github/codeql/tree/main/java/ql/lib/semmle/code/java/security.
  - Github. codeql/java/ql/src/Security/CWE at main · github/codeql, c.
    - Github. execute language-server, d. URL https://docs-internal.github.com/en/code-security/codeql-cli/codeql-cli-manual/execute-language-server.
    - Github. CodeQL. https://codeql.github.com/, 2025.
  - GitHub. Github security advisory database. https://github.com/advisories, 2025.
- Gwihwan Go, Chijin Zhou, Quan Zhang, Yu Jiang, and Zhao Wei. LSPAI: An IDE Plugin for LLM-Powered Multi-Language Unit Test Generation with Language Server Protocol. In *Proceedings* of the 33rd ACM International Conference on the Foundations of Software Engineering, pp. 144–149, Clarion Hotel Trondheim Trondheim Norway, June 2025. ACM. ISBN 979-8-4007-1276-0. doi: 10.1145/3696630.3728540. URL https://dl.acm.org/doi/10.1145/3696630. 3728540.
  - Penghui Li, Songchen Yao, Josef Sarfati Korich, Changhua Luo, Jianjia Yu, Yinzhi Cao, and Junfeng Yang. Automated static vulnerability detection via a holistic neuro-symbolic approach, 2025a. URL https://arxiv.org/abs/2504.16057.
  - Z. Li, S. Dutta, and M. Naik. IRIS: LLM-Assisted Static Analysis for Detecting Security Vulnerabilities, Apr. 2025b. arXiv:2405.17238.
- Meta. Infer. https://fbinfer.com/, 2025.
- MITRE. Common vulnerabilities and exposures (cve). https://www.cve.org/, 2025.
- Federico Mora, Justin Wong, Haley Lepe, Sahil Bhatia, Karim Elmaaroufi, George Varghese, Joseph E Gonzalez, Elizabeth Polgreen, and Sanjit A Seshia. Synthetic programming elicitation for text-to-code in very low-resource programming and formal languages. *Advances in Neural Information Processing Systems*, 37:105151–105170, 2024.
- NIST. National vulnerability database (nvd). https://nvd.nist.gov/, 2025.
- Semgrep. Semgrep. the semgrep platform. https://semgrep.dev/, 2023.
- C. Yang, Z. Zhao, Z. Xie, H. Li, and L. Zhang. KNighter: Transforming Static Analysis with LLM-Synthesized Checkers, Apr. 2025. arXiv:2503.09002.

John Yang, Carlos E. Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. Swe-agent: Agent-computer interfaces enable automated software engineering, 2024. URL https://arxiv.org/abs/2405.15793.

Xin Zhou, Sicong Cao, Xiaobing Sun, and David Lo. Large language model for vulnerability detection and repair: Literature review and the road ahead, 2024. URL https://arxiv.org/abs/2404.02525.

## A CODEQL QUERIES

540

541

542

543

544

545

546547548

549

550 551

552

553

554

555

#### A.1 CODEQL QUERY STRUCTURE TEMPLATE

The template below is given to the LLM agent at the start of the iterative query synthesis task. The prompt instructs the LLM to use the AST nodes, along with the CodeQL LSP and CodeQL references in the vector database, to fill in this template. The prompt also takes note to find similar queries related to the given CVE's vulnerability.

```
556 <sub>1</sub>
      /**
557 <sub>2</sub>
       * @name [Vulnerability Name based on analysis]
        * @description [Description derived from the vulnerability pattern]
        * @problem.severity error
        * @security-severity [score based on severity]
560 5
        * @precision high
561 6
        * @tags security
562 7
563 <sup>8</sup>
        * @kind path-problem
        * @id [unique-id]
564
        */
   10
565 11
      import java
566 <sub>12</sub>
      import semmle.code.java.frameworks.Networking
567 <sub>13</sub>
      import semmle.code.java.dataflow.DataFlow
568 14 import semmle.code.java.dataflow.FlowSources
569 15 import semmle.code.java.dataflow.TaintTracking
       private import semmle.code.java.dataflow.ExternalFlow
570 16
571 17
572 18
      class Source extends DataFlow::Node {
573 <sup>19</sup>
         Source() {
574 <sup>20</sup>
            exists([AST node type from analysis] |
575 <sup>21</sup>
              /st Fill based on AST patterns for sources identified in Phase 1 & 2 st/
              and this.asExpr() = [appropriate mapping]
            )
577 <sub>24</sub>
         }
578 <sub>25</sub>
       }
579 <sub>26</sub>
580 27
      class Sink extends DataFlow::Node {
         Sink() {
581 28
            exists([AST node type] |
582 29
              /* Fill based on AST patterns for sinks */
583 <sup>30</sup>
              and this.asExpr() = [appropriate mapping]
584 <sup>31</sup>
585 32
            ) or
586 <sup>33</sup>
            exists([Alternative AST pattern] |
              /* Additional sink patterns from analysis */
587 34
35
              and [appropriate condition]
588 36
            )
589 <sub>37</sub>
          }
590 <sub>38</sub>
       }
591 39
      class Sanitizer extends DataFlow::Node {
592 40
         Sanitizer() {
593 41
            exists([AST node type for sanitizers] |
   42
```

```
/* Fill based on sanitizer patterns from Phase 1 & 2 */
595 <sub>44</sub>
596 <sub>45</sub>
         }
597 <sub>46</sub>
598 47
      module MyPathConfig implements DataFlow::ConfigSig {
599 48
         predicate isSource(DataFlow::Node source) {
601 50
            source instanceof Source
602 51
603 52
         predicate isSink(DataFlow::Node sink) {
604 53
605
           sink instanceof Sink
   55
606
   56
607 <sub>57</sub>
         predicate isBarrier(DataFlow::Node sanitizer) {
608 <sub>58</sub>
            sanitizer instanceof Sanitizer
609 59
610 60
         predicate isAdditionalFlowStep(DataFlow::Node n1, DataFlow::Node n2) {{
611 61
            /* Fill based on additional taint steps from analysis */
612 62
613 63
      } }
614 64
615 65
616
       module MyPathFlow = TaintTracking::Global<MyPathConfig>;
       import MyPathFlow::PathGraph
617 68
618 <sub>69</sub>
619 <sub>70</sub>
         MyPathFlow::PathNode source,
620 71
         MyPathFlow::PathNode sink
      where
621 72
         MyPathFlow::flowPath(source, sink)
622 73
623 74
      select
         sink.getNode(),
624 75
625 76
         source,
626 77
         "[Alert message based on vulnerability]",
627
         source.getNode(),
628 <sub>80</sub>
         "[source description]"
629
630
```

#### A.2 ITERATIVE FEEDBACK EXAMPLE

631 632

633

635

636

637

The following is an excerpt from the prompt used in a new context window that instructs the agent to improve the last iteration's query. The feedback is derived from automatically running synthesized queries against the fixed and vulnerable versions of the CVE's repository, and comparing the query's results against the changed methods in the CVE fix commit diff. The feedback also includes any successful method hits, false positive method hits, and which fixed methods have not been detected by the query.

```
638
      ## Objective
639
      Refine the CodeQL query based on previous iteration feedback to
   2
640
       → improve vulnerability detection.
641
642 4
      ## Previous Iteration Feedback
643 <sub>5</sub>
      ## Iteration 2 Results
644 6
645 7
      ## Previous Query (Iteration 2)
       ``ql
646 8
647 9
       * @name AntiSamy XSS Bypass in Style Tags
   10
```

```
648
        * @description Detects potential XSS vulnerabilities where
649
        → AntiSamy HTML sanitization fails to properly filter malicious

→ style tags with event handlers

651 12
        * @problem.severity error
652 13
        * @security-severity 6.1
653 14
        * Oprecision high
        * @tags security
654 15
        * @kind path-problem
655 16
        * @id java/antisamy-style-xss-bypass
656 17
        */
657 18
658 <sup>19</sup>
      import java
      import semmle.code.java.dataflow.DataFlow
   20
659
       import semmle.code.java.dataflow.TaintTracking
   21
660 22
661 23
      class MaliciousHtmlInput extends DataFlow::Node {
662 <sub>24</sub>
        MaliciousHtmlInput() {
663 25
           exists(StringLiteral sl |
664 26
             this.asExpr() = sl and
665 27
                sl.getValue().matches("%<style%onload%>%") or
666 28
                sl.getValue().matches("%<style%onclick%>%") or
667 29
                sl.getValue().matches("%<style%onerror%>%") or
668 30
                sl.getValue().matches("%<style%onmouseover%>%") or
669 31
670 32
                sl.getValue().matches("%style%alert%") or
                sl.getValue().matches("%<style%javascript%") or</pre>
671 <sub>34</sub>
                sl.getValue().matches("%<style%expression%")</pre>
672 <sub>35</sub>
             )
673 <sub>36</sub>
           ) or
674 37
           exists(Parameter p |
             this.asParameter() = p and
675 38
676 39
               p.getName().toLowerCase().matches("%html%") or
677 40
                p.getName().toLowerCase().matches("%input%") or
678 41
679 42
               p.getName().toLowerCase().matches("%content%")
680 43
681 44
           )
         }
   45
682 <sub>46</sub>
      }
683 <sub>47</sub>
684 <sub>48</sub>
      class AntiSamyProcessing extends DataFlow::Node {
685 49
         AntiSamyProcessing() {
           exists (MethodCall mc |
686 50
             this.asExpr() = mc and
687 51
             mc.getMethod().hasName("scan") and
688 52
             mc.getMethod().getDeclaringType()(RefType)
689 53
             .hasQualifiedName("org.owasp.validator.html", "AntiSamy")
690 54
691 55
           ) or
           exists(MethodCall mc |
692
             this.asExpr() = mc and
693 <sub>58</sub>
             mc.getMethod().hasName("getCleanHTML") and
694 59
            mc.getQualifier().getType().(RefType)
695 60
             .hasQualifiedName("org.owasp.validator.html",
              696
697 61
698 62
699 63
700 64
701 65
      class PotentialXssOutput extends DataFlow::Node {
         PotentialXssOutput() {
```

```
702 <sub>67</sub>
            exists (MethodCall mc |
703 <sub>68</sub>
              this.asExpr() = mc and
704 <sub>69</sub>
              (
705 <sub>70</sub>
                mc.getMethod().hasName("contains") and
706 71
                 exists(StringLiteral sl |
707 72
                   mc.getAnArgument() = sl and
                   (
708 73
                      sl.getValue().matches("%alert%") or
709 74
                      sl.getValue().matches("%script%") or
710 75
                     sl.qetValue().matches("%javascript%") or
711 76
                     sl.getValue().matches("%onload%") or
712 77
713 78
79
                     sl.getValue().matches("%onclick%")
714 80
                   )
                 )
715 81
              )
716 <sub>82</sub>
            ) or
717 83
            exists(MethodCall mc |
718 84
             this.asExpr() = mc and
              mc.getMethod().hasName("getCleanHTML") and
719 85
             mc.getQualifier().getType().(RefType)
720 86
721 87
              .hasQualifiedName("org.owasp.validator.html",
               → "CleanResults")
722
            )
723 88
          }
724 89
       }
   90
725 91
726 <sub>92</sub>
       class ImproperSanitizationBarrier extends DataFlow::Node {
727 <sub>93</sub>
         ImproperSanitizationBarrier() {
728 94
            exists (MethodCall mc |
729 95
              this.asExpr() = mc and
             mc.getMethod().hasName("assertTrue") and
730 96
             exists(LogNotExpr lne |
731 97
                mc.getAnArgument() = lne and
732 <sup>98</sup>
733 99
                 lne.getExpr().(MethodCall).getMethod().hasName("contains")
734<sup>100</sup>
735 101
            )
          }
736 <sub>103</sub>
       }
737<sub>104</sub>
738<sub>105</sub>
       module AntiSamyXssConfig implements DataFlow::ConfigSig {
739<sub>106</sub>
         predicate isSource(DataFlow::Node source) {
740 107
           source instanceof MaliciousHtmlInput
741 108
742 109
         predicate isSink(DataFlow::Node sink) {
743 110
744 111
            sink instanceof PotentialXssOutput
745
         }
746
   114
         predicate isBarrier(DataFlow::Node sanitizer) {
747
            sanitizer instanceof ImproperSanitizationBarrier
748<sub>116</sub>
749 117
750118
         predicate isAdditionalFlowStep(DataFlow::Node n1, DataFlow::Node
751
          \rightarrow n2) {
            exists (MethodCall mc |
752119
              n1.asExpr() = mc.getQualifier() and
753 120
              n2.asExpr() = mc and
754 121
755 <sup>122</sup>
              mc.getMethod().hasName("scan") and
              mc.getMethod().getDeclaringType().(RefType)
```

```
.hasQualifiedName("org.owasp.validator.html", "AntiSamy")
757
758<sub>126</sub>
           exists(MethodCall mc |
759<sub>127</sub>
             n1.asExpr() = mc.getQualifier() and
760 128
             n2.asExpr() = mc and
             mc.getMethod().hasName("getCleanHTML")
761 129
762130
           )
         }
763 131
      }
764 132
765^{133}
      module AntiSamyXssFlow = TaintTracking::Global<AntiSamyXssConfig>;
767 135
767 136
768 137
      import AntiSamyXssFlow::PathGraph
769<sub>138</sub>
         AntiSamyXssFlow::PathNode source,
770 139
         AntiSamyXssFlow::PathNode sink
771<sub>140</sub> where
772 141
      AntiSamyXssFlow::flowPath(source, sink)
773142 select
      sink.getNode(),
774 143
775 144
        source,
        sink,
776 145
        "Potential XSS vulnerability: HTML input with malicious style
777^{146}

→ tags may bypass AntiSamy sanitization",

778 147
        source.getNode(),
779
        "malicious HTML input"
780 <sub>149</sub>
781 <sub>150</sub>
782 151
      ## Compilation Results
783152 COMPILATION SUCCESS: Query syntax is valid
784153
      ## Execution Results
785 154
      ## Query Evaluation Summary (Iteration 2)
786 155
787<sup>156</sup>
      Results: Vulnerable=8, Fixed=8
788 157
     Method Recall: Vulnerable=True, Fixed=True
789
      True Positive Methods: Vulnerable=2, Fixed=2
      Coverage: 1/1 target methods
790 160 PARTIAL: Query hits targets but has false positives in fixed
791

→ version

792161 Method location format is path/to/hit/file.java:[Class of hit
793
       → method]:[Hit method]
794162 Successfully targeted methods:
795 163
         -src/main/java/org/owasp/validator/html/scan/MagicSAXFilter.java:MagicSAXFilte
      False positives (hits in fixed version):
798 <sup>165</sup>
         → src/main/java/org/owasp/validator/html/scan/MagicSAXFilter.java:MagicSAXFilter
799
800
      ## Detailed Evaluation Analysis
  167
801
802 169
      **Method Coverage**: 1/1 target methods
803<sub>170</sub>
      **File Coverage**: 1/1 target files
804 171
      **Successfully targeted files**:
805 172
          - MagicSAXFilter.java
806173
      **Successfully targeted methods**:
807174
          - startElement
808 175
809 176
      **False positives (hit in fixed version) **:
```

```
- startElement
811
812<sub>180</sub>
      **Evaluation Summary**:
813 181
      - Vulnerable DB File Recall: True
814 182
      - Fixed DB File Recall: True
815<sub>183</sub> - Vulnerable DB Method Recall: True
816184 - Fixed DB Method Recall: True
      - Total Query Results: Vulnerable=8, Fixed=8
817185
      - Code Flow Paths: Vulnerable=8, Fixed=8
818 186
819 187
820<sup>188</sup>
      ## Next Steps
821 189
      We want vulnerable DB method recall and we don t want fixed DB

→ method recall!

822
      **Priority**: Reduce false positives by adding more specific
823
       → conditions to avoid hitting the methods listed above.
824
```

#### A.3 FINENIB SYNTHESIZED QUERY EXAMPLES

825

826 827

828

829

830

831 832 Below are examples of successful synthesized queries. The queries successfully find > 0 true positive method hits on the vulnerable CodeQL database of the CVE's source code, and no false positive method hits on the fixed version's CodeQL database. For reference queries to compare with, CWE queries can be found on the official CodeQL repository (Github, c;b)

CVE-2025-27136, CWE-611 - Improper Restriction of XML External Entity Reference

```
833
      /**
834 1
       * @name XML External Entity vulnerability in WstxInputFactory
        → without secure configuration
836
       * @description WstxInputFactory used in XmlFactory without
837
        → disabling DTD support and external entities allows XXE
838
           attacks
839 4
       * @problem.severity error
840 5
       * @security-severity 9.1
841 6
       * @precision high
842 7
       * @tags security
       * @kind path-problem
843 8
844 9
       * @id java/wstxinputfactory-xxe
845 10
846
847
      import java
13
848
14
      import semmle.code.java.dataflow.DataFlow
      import semmle.code.java.dataflow.FlowSources
849 15
      import semmle.code.java.dataflow.TaintTracking
850 <sub>16</sub>
      private import semmle.code.java.dataflow.ExternalFlow
851 17
852 18
      class WstxInputFactoryCreation extends DataFlow::Node {
        WstxInputFactoryCreation() {
853 19
          exists (MethodCall mc |
854 20
855 21
            mc.getMethod().hasQualifiedName("com.ctc.wstx.stax",
             → "WstxInputFactory", "newInstance") or
856
            mc.getMethod().hasQualifiedName("com.ctc.wstx.stax",
857 22
             → "WstxInputFactory", "newFactory")
858 <sub>23</sub>
859 <sub>24</sub>
             this.asExpr() = mc
860 25
          ) or
861 <sub>26</sub>
          exists(ClassInstanceExpr cie |
862 27
                cie.getConstructedType().hasQualifiedName("com.ctc.wstx.stax",
                 "WstxInputFactory") and
```

```
864 28
            this.asExpr() = cie
865 29
866 30
          // Include variable access to WstxInputFactory instances (like
867
           → "input" parameter)
868 31
          exists (Variable v, VarAccess va |
            v.getType().(RefType).hasQualifiedName("com.ctc.wstx.stax",
869 32
             → "WstxInputFactory") and
            va.getVariable() = v and
871 33
            this.asExpr() = va
872 34
          )
873 35
        }
874 <sup>36</sup>
875 37
876 39
      class UnsafeXmlFactoryUsage extends DataFlow::Node {
877 40
        UnsafeXmlFactoryUsage() {
878 <sub>41</sub>
          exists(ClassInstanceExpr xmlFactoryCall |
879 42
             // XmlFactory constructor with WstxInputFactory parameter
880 43
            xmlFactoryCall.getConstructedType()
             .hasQualifiedName("com.fasterxml.jackson.dataformat.xml",
881 44

→ "XmlFactory") and

882
883 45
            xmlFactoryCall.getArgument(0) = this.asExpr()
          ) or
884 46
          exists(ClassInstanceExpr xmlMapperCall, ClassInstanceExpr
885 47
           886
            // XmlMapper constructor using XmlFactory with
887
             → WstxInputFactory
888 49
            xmlMapperCall.getConstructedType()
889 <sub>50</sub>
             .hasQualifiedName("com.fasterxml.jackson.dataformat.xml",
890

→ "XmlMapper") and

            xmlFactoryCall.getConstructedType()
891 51
             .hasQualifiedName("com.fasterxml.jackson.dataformat.xml",
892 52
             893
            xmlMapperCall.getArgument(0) = xmlFactoryCall and
894 53
            xmlFactoryCall.getArgument(0) = this.asExpr()
895 54
896 55
          )
897 56
        }
898 58
899 59
      class WstxInputFactorySanitizer extends DataFlow::Node {
900 60
        WstxInputFactorySanitizer() {
901 61
          exists (MethodCall setPropertyCall, VarAccess factoryVar |
902 62

→ setPropertyCall.getMethod().hasQualifiedName("javax.xml.stream",
903
                 "XMLInputFactory", "setProperty") and
             setPropertyCall.getQualifier() = factoryVar and
905 63
906 64
907 65
               // DTD support disabled
               (exists(Field f |
908
                 setPropertyCall.getArgument(0) = f.getAnAccess() and
909 68
                 f.hasName("SUPPORT_DTD") and
910 69
911
                     f.getDeclaringType().hasQualifiedName("javax.xml.stream",
912
                     "XMLInputFactory")
                 \hookrightarrow
913 70
               ) and
               exists(Field f |
914 71
                 setPropertyCall.getArgument(1) = f.getAnAccess() and
915 72
                 f.hasName("FALSE") and
916 73
                 f.getDeclaringType().hasQualifiedName("java.lang",
917 74
                 → "Boolean")
```

```
918 75
               )) or
919 <sub>76</sub>
               // External entities disabled
920 77
               (exists (Field f |
921 78
                 setPropertyCall.getArgument(0) = f.getAnAccess() and
922 79
                 f.hasName("IS_SUPPORTING_EXTERNAL_ENTITIES") and
923 80
                     f.getDeclaringType().hasQualifiedName("javax.xml.stream",
924
                  → "XMLInputFactory")
925
               ) and
926 81
               exists (Field f |
927 82
                 setPropertyCall.getArgument(1) = f.getAnAccess() and
928 83
                 f.hasName("FALSE") and
929
                 f.getDeclaringType().hasQualifiedName("java.lang",
930
                  → "Boolean")
931 86
               ))
932 87
             ) and
933 88
             this.asExpr() = factoryVar
934 89
          )
        }
935 90
      }
936 91
937 92
      module WstxInputFactoryFlowConfig implements DataFlow::ConfigSig {
938 93
        predicate isSource(DataFlow::Node source) {
939 94
          source instanceof WstxInputFactoryCreation
940 95
941 97
942 98
        predicate isSink(DataFlow::Node sink) {
943 99
          sink instanceof UnsafeXmlFactoryUsage
944_{100}
945 101
        predicate isBarrier(DataFlow::Node sanitizer) {
946102
          sanitizer instanceof WstxInputFactorySanitizer
947103
948 104
949 105
950 106
        predicate isAdditionalFlowStep(DataFlow::Node n1, DataFlow::Node
         \rightarrow n2) {
951
          // WstxInputFactory passed to XmlFactory constructor
952
          exists(ClassInstanceExpr xmlFactoryCall |
953 109
             xmlFactoryCall.getConstructedType()
954 110
             .hasQualifiedName("com.fasterxml.jackson.dataformat.xml",
955

→ "XmlFactory") and

             xmlFactoryCall.getArgument(0) = n1.asExpr() and
956111
             n2.asExpr() = xmlFactoryCall
957112
          ) or
958113
           // XmlFactory passed to XmlMapper constructor
959 114
          exists(ClassInstanceExpr xmlMapperCall |
960 115
961 116
             xmlMapperCall.getConstructedType()
             .hasQualifiedName("com.fasterxml.jackson.dataformat.xml",
             963
             xmlMapperCall.getArgument(0) = n1.asExpr() and
964 119
             n2.asExpr() = xmlMapperCall
965 120
          )
966 121
        }
967122
968 123
      module WstxInputFactoryFlow =
969 124
      → TaintTracking::Global<WstxInputFactoryFlowConfig>;
970
971 125
      import WstxInputFactoryFlow::PathGraph
  126
```

```
from
973
        WstxInputFactoryFlow::PathNode source,
974 129
        WstxInputFactoryFlow::PathNode sink
975<sub>130</sub>
      where
976 131
      WstxInputFactoryFlow::flowPath(source, sink)
977<sub>132</sub> select
      sink.getNode(),
978 133
      source,
979 134
        sink,
980 135
        "WstxInputFactory used without secure configuration flows to XML
981 136
        → parser, allowing XXE attacks",
982
983
        source.getNode(),
138
984
        "WstxInputFactory usage"
985
      CVE-2025-0851, CWE-22 - Path Traversal
986
988 2
      * @name Archive path traversal vulnerability (ZipSlip) -
       → CVE-2025-0851
990 3
       * @description Archive entries with path traversal sequences can
        → write files outside the intended extraction directory
       * @problem.severity error
992 4
       * @security-severity 9.8
993 5
       * @precision high
       * @tags security
995 7
996 8
      * @kind path-problem
       * @id java/archive-path-traversal-cve-2025-0851
997
998 11
999_{12} import java
1000<sub>13</sub> import semmle.code.java.dataflow.DataFlow
100114 import semmle.code.java.dataflow.TaintTracking
100215
     /**
100316
100417
       * Sources: Archive entry names from ZipEntry.getName() and
       → TarArchiveEntry.getName()
1005
1006
1007
     class ArchiveEntryNameSource extends DataFlow::Node {
      ArchiveEntryNameSource() {
1008 21
          exists (MethodCall mc |
1009<sub>22</sub>
            mc.getMethod().getName() = "getName" and
1010
             → (mc.getMethod().getDeclaringType().hasQualifiedName("java.util.zip",
1011
             101223
                   .getDeclaringType()
                    → .hasQualifiedName("org.apache.commons.compress.archivers.tar",
1013
                      "TarArchiveEntry")
1014
             ) and
1015^{24}
             this.asExpr() = mc
1016<sup>25</sup>
1017<sup>26</sup>
101827
        }
1019 28
      }
102030
102131
     * Sinks: Path resolution operations that lead to file creation
102232
102333 class PathCreationSink extends DataFlow::Node {
        PathCreationSink() {
102434
          // Arguments to Path.resolve() calls
102535
          exists(MethodCall resolveCall |
  36
```

```
1026
37
             resolveCall.getMethod().getName() = "resolve" and
1027
             resolveCall.getMethod().getDeclaringType()
102839
              .hasQualifiedName("java.nio.file", "Path") and
1029_{40}
             this.asExpr() = resolveCall.getAnArgument()
1030_{41}
           )
103142
           or
           // Arguments to file creation operations
103243
           exists (MethodCall fileOp |
103344
             (
1034^{45}
                fileOp.getMethod().getName() = "createDirectories" or
1035<sup>46</sup>
                fileOp.getMethod().getName() = "newOutputStream" or
1036
1037
                fileOp.getMethod().getName() = "write" or
                fileOp.getMethod().getName() = "copy"
1038
             ) and
1039<sub>51</sub>
1040
              → fileOp.getMethod().getDeclaringType().hasQualifiedName("java.nio.file",
1041
                 "Files") and
104252
             this.asExpr() = fileOp.getAnArgument()
104353
         }
104454
      }
1045<sup>55</sup>
1046<sup>56</sup>
1047<sup>57</sup>
1048.58
      * Sanitizers: Proper validation that prevents path traversal
104960
      class PathTraversalSanitizer extends DataFlow::Node {
1050<sub>61</sub>
         PathTraversalSanitizer() {
105162
           // The validateArchiveEntry method call that properly
1052
               validates paths
           // This blocks flow after the validation call is made
105363
           exists(MethodCall validateCall |
105464
             validateCall.getMethod().getName() = "validateArchiveEntry"
105565
              → and
1056
              (
1057<sup>66</sup>
                // Any variable assigned from validateArchiveEntry call
1058<sup>67</sup>
                → result
1059
                exists(Variable v |
1060<sub>69</sub>
                 this.asExpr() = v.getAnAccess() and
1061<sub>70</sub>
                  exists(AssignExpr assign |
1062_{71}
                     assign.getDest() = v.getAnAccess() and
106372
                     assign.getRhs() = validateCall
                  )
106473
                )
106574
                or
106675
                // Variables passed through validateArchiveEntry calls
106776
                this.asExpr() = validateCall.getAnArgument() and
1068<sup>77</sup>
                exists(ExprStmt stmt | stmt.getExpr() = validateCall)
1069<sup>78</sup>
             )
1070
           )
1071<sub>81</sub>
           or
1072<sub>82</sub>
           // Proper ".." validation with exception throwing (complete
1073
           → pattern)
107483
           exists (MethodCall containsCall, IfStmt ifStmt, ThrowStmt
1075

    → throwStmt |

             containsCall.getMethod().getName() = "contains" and
107684
             containsCall.getAnArgument().(StringLiteral).getValue() =
107785

→ ".." and

1078
             ifStmt.getCondition().getAChildExpr*() = containsCall and
107986
             ifStmt.getThen().getAChild*() = throwStmt and
```

```
1080
             this.asExpr() = containsCall.getQualifier()
108189
           )
108290
           or
108391
           // Path normalization combined with startsWith validation
108492
           exists(MethodCall normalizeCall, MethodCall startsWithCall |
             normalizeCall.getMethod().getName() = "normalize" and
108593
             normalizeCall.getMethod().getDeclaringType()
108694
             .hasQualifiedName("java.nio.file", "Path") and
108795
             startsWithCall.getMethod().getName() = "startsWith" and
1088<sup>96</sup>
             startsWithCall.getMethod().getDeclaringType().
108997
             hasQualifiedName("java.nio.file", "Path") and
109098
             DataFlow::localFlow(DataFlow::exprNode(normalizeCall),
1091
               DataFlow::exprNode(startsWithCall.getQualifier())) and
1092
             this.asExpr() = normalizeCall.getQualifier()
1093
109403
109 $\overline{4}04
      }
109605
109706
       * Additional predicate to detect validation barriers at method
109407
1099
       */
110^{108}
      predicate hasValidationCall(Callable method) {
110109
1102
        exists (MethodCall validateCall |
           validateCall.getEnclosingCallable() = method and
  111
1103
112
           validateCall.getMethod().getName() = "validateArchiveEntry"
1104
113
         )
110514
110615
      module PathTraversalConfig implements DataFlow::ConfigSig {
110716
        predicate isSource(DataFlow::Node source) {
110817
           source instanceof ArchiveEntryNameSource
110918
1110^{19}
1111^{120}
1112
        predicate isSink(DataFlow::Node sink) {
1112
1113
123
           sink instanceof PathCreationSink
111<u>4</u>
124
1115_{25}
        predicate isBarrier(DataFlow::Node sanitizer) {
111626
           sanitizer instanceof PathTraversalSanitizer
111727
111828
        predicate isBarrierIn(DataFlow::Node node) {
111 129
          // Barrier at method entry if method contains
112030
           → validateArchiveEntry call
1121
           node instanceof DataFlow::ParameterNode and
112231
1123^{\overset{\frown}{1}32}
           hasValidationCall(node.getEnclosingCallable())
133
1124
134
1125
135
        predicate isAdditionalFlowStep(DataFlow::Node n1, DataFlow::Node
1126
         \rightarrow n2) {
112736
           // Flow through variable assignments and declarations
112937
           exists(LocalVariableDeclExpr decl |
112938
             decl.getInit() = n1.asExpr() and
             n2.asExpr() = decl.getVariable().getAnAccess()
113139
           )
113140
           or
113^{41}
113342
           exists(AssignExpr assign |
             assign.getRhs() = n1.asExpr() and
  143
```

```
1134
144
            n2.asExpr() = assign.getDest()
1135
145
          )
113646
          or
113747
          // Flow through string manipulation methods that preserve
1138

    → taint

113948
          exists(MethodCall mc |
          mc.getAnArgument() = n1.asExpr() and
114049
           n2.asExpr() = mc and
114150
114251
              mc.getMethod().getName() = "removeLeadingFileSeparator" or
114352
1144^{\overset{153}{4}}
              mc.getMethod().getName() = "trim" or
1145
              mc.getMethod().getName() = "toString" or
              mc.getMethod().getName() = "substring"
1146
156
            )
1147
          )
114858
          or
114959
          // Flow through Path operations
          exists(MethodCall pathOp |
115060
           pathOp.getAnArgument() = n1.asExpr() and
115161
           n2.asExpr() = pathOp and
115262
            pathOp.getMethod().getName() = "resolve" and
1153^{63}
            pathOp.getMethod().getDeclaringType()
1154^{64}
            .hasQualifiedName("java.nio.file", "Path")
1155^{165}
1156
          )
  167
        }
1157
      }
1158
1159<sub>70</sub>
      module PathTraversalFlow =
1160
      → TaintTracking::Global<PathTraversalConfig>;
116171
116172 import PathTraversalFlow::PathGraph
116373
     from PathTraversalFlow::PathNode source,
116474
      → PathTraversalFlow::PathNode sink
1165
1166
1167
        PathTraversalFlow::flowPath(source, sink) and
        // Focus on the specific vulnerable files and methods
1168
178
116979
          source.getNode().getEnclosingCallable().getDeclaringType()
117 \rho_{80}
          .hasName("TarUtils") or
117181
          source.getNode().getEnclosingCallable().getDeclaringType()
          .hasName("ZipUtils")
117282
      ) and
117183
117484
          source.getNode().getEnclosingCallable().getName() = "untar" or
1175^{85}
1176^{186}
          source.getNode().getEnclosingCallable().getName() = "unzip"
1177
        ) and
        // Only report flows where validation is NOT properly done
1178
        not hasValidationCall(source.getNode().getEnclosingCallable())
1179 select sink.getNode(), source, sink,
1180
        "Archive entry name from $@ flows to file system operation
1181
        1182
        → attack."
        source.getNode(), "archive entry name"
118392
1184
      CCVE-2025-27528, CWE-502 - Deserialization of Untrusted Data
1185
1186
11871 /**
      * @name MySQL JDBC URL parameter injection vulnerability
```

```
1188 3
        * @description Detects MySQL JDBC URLs with dangerous bracket
1189
        → parameters that bypass inadequate filtering in vulnerable
1190
        → code
1191 4
        * @problem.severity error
1192 5
        * @security-severity 8.8
        * @precision high
1193 6
        * @tags security
11947
        * @kind path-problem
11958
        * @id java/mysql-jdbc-url-injection
1196 <sup>9</sup>
1197<sup>10</sup>
1198
1199
       import java
       import semmle.code.java.dataflow.DataFlow
1200
       import semmle.code.java.dataflow.TaintTracking
1201<sub>15</sub>
1202<sub>16</sub>
      class MySQLDangerousBracketUrlSource extends DataFlow::Node {
1203<sub>17</sub>
         MySQLDangerousBracketUrlSource() {
120418
           // String literals with dangerous MySQL parameters in bracket
            → notation
1205
           exists(StringLiteral lit |
120619
             lit.getValue().matches("*mysql*") and
1207<sup>20</sup>
             lit.getValue().matches("*[*]*") and
1208<sup>21</sup>
1209<sup>22</sup>
1210 23
                lit.getValue().matches("*allowLoadLocalInfile*") or
                lit.getValue().matches("*allowUrlInLocalInfile*") or
1211<sub>25</sub>
                lit.getValue().matches("*autoDeserialize*") or
1212
                lit.getValue().matches("*allowPublicKeyRetrieval*") or
1213<sub>27</sub>
                lit.getValue().matches("*serverTimezone*") or
1214<sub>28</sub>
                lit.getValue().matches("*user*") or
                lit.getValue().matches("*password*")
121529
             ) and
121630
             this.asExpr() = lit
121731
           )
1218<sup>32</sup>
           or
1219<sup>33</sup>
           // Parameters to filterSensitive method that may contain
122034
            → dangerous bracket content
1221<sub>35</sub>
           exists (Method m, Parameter p |
1222
             m.hasName("filterSensitive") and
1223<sub>37</sub>
             m.getDeclaringType().getName() = "MySQLSensitiveUrlUtils"
1224

→ and

122538
             p = m.getAParameter() and
             this.asParameter() = p
122639
122740
122841
      }
122942
1230<sup>43</sup>
       class VulnerableCodePatternSink extends DataFlow::Node {
1231<sup>44</sup>
         VulnerableCodePatternSink() {
1232
           // The vulnerability: calls to filterSensitive in vulnerable
1233
            → code patterns
123447
           exists(Method m, MethodCall filterCall |
123548
             m.hasName("filterSensitive") and
             m.getDeclaringType().getName() = "MySQLSensitiveUrlUtils"
123640
1237
             filterCall.getMethod() = m and
123850
             this.asExpr() = filterCall and
123951
             // Key vulnerability condition: this code exists where
1240<sup>52</sup>
              → filterSensitiveKeyByBracket method is NOT available
1241
```

```
1242
53
             // In the vulnerable version, filterSensitiveKeyByBracket
1243
              → doesn't exist
1244<sub>54</sub>
             not exists (Method bracketMethod |
124555
                bracketMethod.hasName("filterSensitiveKeyByBracket") and
1246<sub>56</sub>
                bracketMethod.getDeclaringType().getName() =
                → "MySQLSensitiveUrlUtils" and
1247
                bracketMethod.getDeclaringType() = m.getDeclaringType()
124857
             )
124958
           )
1250<sup>59</sup>
           or
1251<sup>60</sup>
           // Additional sink: method calls that use the result of
125261
            → inadequate filtering
1253<sub>62</sub>
           exists (MethodCall mc, MethodCall filterCall |
1254<sub>63</sub>
             filterCall.getMethod().hasName("filterSensitive") and
1255<sub>64</sub>
             filterCall.getMethod().getDeclaringType().getName() =
1256
              → "MySQLSensitiveUrlUtils" and
125765
             DataFlow::localFlow(DataFlow::exprNode(filterCall),
1258
              → DataFlow::exprNode(mc.getArgument(_))) and
             this.asExpr() = mc and
125966
             // Only vulnerable if no proper bracket filtering exists in
126067
             \rightarrow the same class
1261
             not exists(Method bracketMethod |
1262<sup>68</sup>
               bracketMethod.hasName("filterSensitiveKeyByBracket") and
1263<sup>69</sup>
                bracketMethod.getDeclaringType().getName() =
1264
                → "MySQLSensitiveUrlUtils" and
1265<sub>71</sub>
                bracketMethod.getDeclaringType() =
1266
                → filterCall.getMethod().getDeclaringType()
126772
126873
126974
127075
127176
      class ProperBracketFilteringSanitizer extends DataFlow::Node {
127277
1273<sup>78</sup>
      ProperBracketFilteringSanitizer() {
           // The proper bracket-based sanitization method (present only
1274

→ in fixed version)

1275
           exists (MethodCall mc |
127681
             mc.getMethod().hasName("filterSensitiveKeyByBracket") and
1277<sub>82</sub>
             mc.getMethod().getDeclaringType().getName() =
              → "MySQLSensitiveUrlUtils" and
127983
             this.asExpr() = mc
           )
128084
         }
128185
      }
128286
128387
      module MySQLJDBCUrlInjectionConfig implements DataFlow::ConfigSig
1284<sup>88</sup>
1285
        predicate isSource(DataFlow::Node source) {
1286 90
           source instanceof MySQLDangerousBracketUrlSource
1287
1288<sub>92</sub>
1289<sub>93</sub>
         predicate isSink(DataFlow::Node sink) {
129001
           sink instanceof VulnerableCodePatternSink
129195
129296
         predicate isBarrier(DataFlow::Node sanitizer) {
1293<sup>97</sup>
           sanitizer instanceof ProperBracketFilteringSanitizer
1294<sup>98</sup>
1295<sup>99</sup>
  100
```

```
1296
101
         predicate isAdditionalFlowStep (DataFlow::Node n1, DataFlow::Node
1297
1298
           // Flow through string concatenation operations
1299_{03}
           exists(AddExpr addExpr |
130004
              n1.asExpr() = addExpr.getLeftOperand() and
130105
              n2.asExpr() = addExpr
           )
130106
           or
130307
           exists(AddExpr addExpr |
1304^{08}
              n1.asExpr() = addExpr.getRightOperand() and
1305 109
1306<sup>110</sup>
              n2.asExpr() = addExpr
1307
           )
   112
           or
1308
113
           // Flow through variable assignments
1309
           exists (Assignment assign |
131Q<sub>15</sub>
              n1.asExpr() = assign.getSource() and
131116
              n2.asExpr() = assign.getDest()
131217
           )
131318
           or
           // Flow through return statements
131419
           exists(ReturnStmt ret |
1315^{20}
             n1.asExpr() = ret.getResult() and
1316<sup>21</sup>
              n2.asParameter() =
131722
              → ret.getEnclosingCallable().getAParameter()
1318
123
1319
124
         }
1320
       }
132126
132727
      module MySQLJDBCUrlInjectionFlow =
       → TaintTracking::Global<MySQLJDBCUrlInjectionConfig>;
1323
132428
       import MySQLJDBCUrlInjectionFlow::PathGraph
1325^{29}
132<sup>§30</sup>
       from MySQLJDBCUrlInjectionFlow::PathNode source,
132731
       → MySQLJDBCUrlInjectionFlow::PathNode sink
1328
1320
1329
133
      where MySQLJDBCUrlInjectionFlow::flowPath(source, sink)
       select sink.getNode(), source, sink,
1330
I34
         "MySQL JDBC URL with dangerous bracket parameters flows to
1331
          → vulnerable filtering logic at $@ that lacks proper
1332
          \rightarrow bracket-based sanitization",
133\mathfrak{P}_{35}
         source.getNode(), "dangerous URL source"
1334
       A.4 AST EXTRACTION QUERY
1335
1336
       Given a fix diff, FineNib automatically parses the changed methods and files, and inserts them into
1337
       an AST pretty printing query template. Below is an example of the AST extraction query used for
1338
      CVE-2014-7816.
1339
1340
1341
        * @name Expressions and statements for CVE-2014-7816 changed code areas
1342
        * @description Extract expressions and statements from vulnerability fix areas
1343
        * @id java/expr-stmt-diff-CVE_2014_7816
1344
        * @kind problem
1345
        * @problem.severity recommendation
1346
        */
1347
       import java
1348
1349
       from Element e, Location 1
```

```
1350
      where
1351
        1 = e.getLocation() and ((1.getFile().getBaseName() = "PathSeparatorHandler.java"
1352
        and l.getStartLine() >= 1 and l.getEndLine() <= 100) or
1353
        (l.getFile().getBaseName() = "URLDecodingHandler.java"
        and l.getStartLine() >= 17 and l.getEndLine() <= 128)</pre>
1356
        or (l.getFile().getBaseName() = "ResourceHandler.java"
        and l.getStartLine() >= 158 and l.getEndLine() <= 172)
1358
1359
        or (1.getFile().getBaseName() = "io.undertow.server.handlers.builder.HandlerBuilde
1360
        and l.getEndLine() <= 29)</pre>
1361
1362
        or (l.getFile().getBaseName() = "DefaultServlet.java"
1363
        and l.getStartLine() >= 39
1364
        and l.getEndLine() <= 150)
1365
        or (l.getFile().getBaseName() = "ServletPathMatches.java"
1366
        and l.getStartLine() >= 32
1367
        and l.getEndLine() <= 140))
1368
      select e,
1369
        e.toString() as element,
1370
        e.getAPrimaryQlClass() as elementType,
1371
        l.getFile().getBaseName() as file,
        l.getStartLine() as startLine,
1373
        l.getEndLine() as endLine,
1374
        1.getStartColumn() as startColumn,
1375
        1.getEndColumn() as endColumn
1376
         CODEQL LANGUAGE SERVER VIA MCP
1378
```

The following are the MCP tool specifications and example usage for our custom CodeQL LSP client, wrapped as an MCP server.

#### **B.1** TOOL SPECIFICATIONS

**codeql\_complete** Provides code completions at a specific position in a CodeQL file. Supports pagination for large completion lists and trigger character-based completions.

#### **Inputs:**

1379

1380

1381

1382 1383

1384

1385

1386 1387

1388

1390

1392

1393 1394

1395 1396

1397

- file\_uri (string): The URI of the CodeQL file
- line (number): Line number (0-based)
- character (number): Character position in the line (0-based)
- trigger\_character (string, optional): Optional trigger character (e.g., ".", "::")
- limit (number, optional): Maximum number of completion items to return (default: 50)
- offset (number, optional): Starting position for pagination (default: 0)

**Returns:** CompletionList with pagination metadata containing completion items, each with label, kind, documentation, and text edit information.

Example usage:

```
1400 {
1401  "tool": "codeql_complete",
1402  "arguments": {
1403  "file_uri": "file:///workspace/security-query.ql",
1403  "line": 5,
```

```
1404
             "character": 12,
1405
             "trigger_character": ".",
1406
             "limit": 25
1407
          }
1408
       }
1409
       codeq1_hover Retrieves hover information (documentation, type information) at a specific po-
1410
       sition. Provides rich markdown documentation for CodeQL predicates, classes, and modules.
1411
1412
1413
             • file_uri (string): The URI of the CodeQL file
1414
1415
             • line (number): Line number (0-based)
1416
             • character (number): Character position in the line (0-based)
1417
1418
       Returns: Hover | null containing documentation content in markdown or plain text format,
1419
       with optional range highlighting.
1420
       Example usage:
1421
1422
1423
          "tool": "codeql_hover",
          "arguments": {
1424
             "file_uri": "file:///workspace/security-query.ql",
1425
             "line": 8,
1426
             "character": 15
1427
          }
1428
       }
1429
1430
       codeql_definition Navigates to the definition location for a symbol at a specific position.
1431
       Supports both single definitions and multiple definition locations.
1432
       Inputs:
1433
1434
             • file_uri (string): The URI of the CodeQL file
1435
             • line (number): Line number (0-based)
1437
             • character (number): Character position in the line (0-based)
1438
       Returns: Location | Location[] | null containing URI and range information for def-
1439
       inition locations.
1440
1441
       Example usage:
1442
1443
          "tool": "codeql_definition",
1444
          "arguments": {
1445
             "file_uri": "file:///workspace/security-query.ql",
1446
             "line": 12,
1447
             "character": 8
1448
          }
1449
       }
1450
1451
       codeql_references Finds all references to a symbol at a specific position across the
1452
       workspace. Includes both usage references and declaration references.
1453
       Inputs:
1454
```

• line (number): Line number (0-based)

• file\_uri (string): The URI of the CodeQL file

• character (number): Character position in the line (0-based)

1455

1456

1457

```
1458
       Returns: Location[] | null containing an array of all reference locations with URI and
1459
       range information.
1460
       Example usage:
1461
1462
1463
          "tool": "codeql_references",
1464
          "arguments": {
1465
            "file_uri": "file:///workspace/security-query.ql",
            "line": 6,
1466
             "character": 20
1467
          }
1468
       }
1469
1470
       codeq1_diagnostics Retrieves diagnostics (errors, warnings, information messages) for a
1471
       CodeQL file. Provides real-time syntax and semantic analysis results.
1472
       Inputs:
1473
1474
             • file_uri (string): The URI of the CodeQL file
1475
1476
       Returns: Diagnostic[] containing an array of diagnostic objects with severity, message, range,
1477
       and optional related information.
1478
       Example usage:
1479
1480
1481
          "tool": "codeql_diagnostics",
1482
          "arguments": {
1483
             "file_uri": "file:///workspace/security-query.ql"
1484
       }
1485
1486
       codeql_format Formats a CodeQL file or a specific selection within the file according to Cod-
1487
       eQL style guidelines.
1488
       Inputs:
1489
1490
             • file_uri (string): The URI of the CodeQL file
1491
1492
             • range (Range, optional): Optional range to format with start and end positions
1493
       Returns: TextEdit[] containing an array of text edits that describe the formatting changes to be
1494
       applied.
1495
       Example usage:
1496
1497
1498
          "tool": "codeql format",
1499
          "arguments": {
1500
            "file_uri": "file:///workspace/security-query.ql",
1501
             "range": {
1502
               "start": { "line": 10, "character": 0 },
1503
               "end": { "line": 25, "character": 0 }
1504
             }
```

**codeql\_update\_file** Updates the content of an open CodeQL file in the language server. This allows for dynamic content modification and analysis of unsaved changes.

#### **Inputs:**

}

1505

1506 1507

1508

1509 1510

1511

• file\_uri (string): The URI of the CodeQL file

• content (string): The new complete content of the file **Returns:** string containing a success confirmation message. Example usage: "tool": "codeql\_update\_file", "arguments": { "file\_uri": "file:///workspace/security-query.ql", "content": "import cpp\n\nfrom Function f\nwhere f.hasName(\"strcpy\")\nselect f } 

#### C EVALUATION DETAILS

Table 6 is a more detailed breakdown of the successful query synthesis rate by CWE.

#### C.1 EVALUATION LIMITATIONS

**Codex CLI.** Claude Code and Gemini CLI allow users to configure how many max turns an agent can take in a context window. As of 9/23/2025, Codex CLI does not offer this configuration. Thus we were not able to force Codex to always take up to 50 max turns each context window.

**IRIS.** The original IRIS evaluation consists of 120 Java projects from CWE-Bench-Java. Many of these projects are old with deprecated dependencies, thus we were only able to build and use 112 of the projects with CodeQL 2.22.2. As of 9/23/2025, IRIS supports 11 CWEs and out of the 65 2025 CVEs, we were able to use 24 of them with IRIS. When running some of the IRIS queries, the amount of sources and sink predicates in the query led to out of memory errors. This impacted 9 out of the 24 IRIS 2025 CVE queries, thus we treat those as queries with 0 results and false recall.

Table 6: FineNib Query Success by CWE Type

CWE Type	<b>Total CVEs</b>	# Success	Success (%)	Avg Precision
CWE-022 (Path Traversal)	48	31	64.6%	0.750
CWE-079 (Cross-Site Scripting)	36	18	50.0%	0.621
CWE-094 (Code Injection)	20	12	60.0%	0.606
CWE-078 (OS Command Injection)	12	7	58.3%	0.628
CWE-502 (Deserialization)	6	4	66.7%	0.853
CWE-611 (XXE)	5	3	60.0%	0.657
CWE-287 (Authentication)	4	1	25.0%	0.875
CWE-200 (Information Exposure)	3	0	0.0%	0.667
CWE-400 (Resource Consumption)	3	1	33.3%	0.556
CWE-532 (Information Exposure)	3	3	100.0%	0.686
CWE-770 (Resource Exhaustion)	3	1	33.3%	0.444
CWE-020 (Improper Input Validation)	2	2	100.0%	0.650
CWE-089 (SQL Injection)	2	2	100.0%	1.000
CWE-1333 (ReDoS)	2	0	0.0%	0.000
CWE-284 (Access Control)	2	0	0.0%	0.500
CWE-862 (Authorization)	2	0	0.0%	0.000
CWE-918 (SSRF)	2	1	50.0%	0.500
CWE-023 (Relative Path Traversal)	1	1	100.0%	1.000
CWE-044 (Path Equivalence)	1	1	100.0%	0.667
CWE-083 (Improper Neutralization)	1	1	100.0%	0.052
CWE-1325 (Improperly Controlled Memory)	1	0	0.0%	0.000
CWE-164 (Foreign Code)	1	0	0.0%	0.000
CWE-178 (Case Sensitivity)	1	0	0.0%	1.000
CWE-190 (Integer Overflow)	1	0	0.0%	0.000
CWE-264 (Permissions)	1	0	0.0%	0.000
CWE-267 (Privilege Defined)	1	0	0.0%	0.000
CWE-276 (Incorrect Permissions)	1	0	0.0%	1.000
CWE-285 (Improper Authorization)	1	1	100.0%	1.000
CWE-288 (Authentication Bypass)	1	0	0.0%	0.000
CWE-290 (Authentication Bypass)	1	1	100.0%	1.000
CWE-297 (Improper Certificate)	1	1	100.0%	1.000
CWE-312 (Cleartext Storage)	1	0	0.0%	0.000
CWE-327 (Cryptographic Issues)	1	0	0.0%	0.000
CWE-346 (Origin Validation)	1	0	0.0%	0.200
CWE-352 (CSRF)	1	1	100.0%	0.941
CWE-426 (Untrusted Search Path)	1	0	0.0%	0.000
CWE-835 (Infinite Loop)	1	0	0.0%	0.000
CWE-863 (Authorization)	1	1	100.0%	1.000
Total	176	94	53.4%	0.631