

Developer-Centric Bug Prediction: A Local Desktop Vulnerability Analysis System with Dual-Engine Detection, SQLite Persistence, and Longitudinal Security Dashboards

Mr. B. Srinivasa Rao¹, S. Tejasri², D. Mounika³, K. Gayathri⁴, M. Lokesh⁵

¹Assistant Professor, Department of AI&DS, Sai Spurthi Institute of Technology, Sathupally, B. Gangaram, Telangana, India

^{2,3,4,5}Student, Department of AI&DS, Sai Spurthi Institute of Technology, Sathupally, B. Gangaram, Telangana, India

Abstract—Software vulnerabilities discovered late in the development lifecycle impose disproportionate remediation costs, with post-release fixes reported to be up to 100 times more expensive than those caught at the coding stage. Yet developer-accessible, lightweight tooling that simultaneously provides local pattern-based detection, historical trend visibility, and persistent scan records remains largely absent from the open-source ecosystem. This paper presents Developer-Centric Bug Prediction (DCBP), a cross-platform Python-Tkinter desktop application that integrates a dual-engine static analysis pipeline—a built-in regular-expression vulnerability scanner and an optional Semgrep subprocess integration—with an SQLite-backed persistence layer and a tabbed dashboard-history interface. The built-in engine matches source code line-by-line against 35 curated patterns spanning nine OWASP-aligned vulnerability classes including SQL injection, cross-site scripting, command injection, path traversal, insecure deserialization, weak cryptography, hardcoded credentials, information disclosure, and XML external entity injection. Findings from both engines are normalised into a uniform schema, severity-classified into five levels (Critical, High, Medium, Low, Info), persisted with full metadata and JSON findings snapshots, and rendered with colour-coded severity tags. The dashboard aggregates cumulative statistics across all stored scans; the history tab provides drilldown to individual scan records. Export functionality produces JSON, CSV, and HTML artefacts. Evaluated against a benchmark of 40 deliberately vulnerable Python modules covering all nine vulnerability classes, the built-in engine achieves a detection precision of 91.3%, recall of 87.6%, and F1-score of 0.894. Semgrep augmentation increases recall by 9.2 percentage points on average across classes. The system passes 100% of 48 functional, integration, and performance test cases. Comparative analysis against SonarQube, Semgrep CLI, Bandit, and OWASP Dependency-Check confirms that DCBP uniquely combines local-first operation, zero-configuration persistence, longitudinal dashboards, and multi-format export in a single, zero-cost application.

Keywords—*Static Application Security Testing, Vulnerability Detection, Developer Security Tools, Python Tkinter, SQLite Persistence, Semgrep Integration, OWASP, Secure Coding, Bug Prediction, DevSecOps*

I. INTRODUCTION

Modern software delivery pipelines compress release cycles through agile and DevOps methodologies, creating conditions under which security defects propagate at a rate that outpaces manual review. Static Application Security Testing (SAST) represents the foundational defensive layer in the "shift-left" security paradigm, enabling developers to detect vulnerability patterns in source code before execution, integration, or deployment [1]. The economic case for early

detection is compelling: empirical studies consistently demonstrate that remediation costs increase by an order of magnitude with each lifecycle stage, reaching a 100:1 ratio between post-release patching and design-phase correction [2].

Despite broad availability of SAST tooling, adoption rates among individual developers and small teams remain low. Enterprise platforms such as SonarQube [3] require dedicated server infrastructure, database administration, and licence fees that are prohibitive for academic learners and startup teams. Lightweight alternatives such as Bandit [4] and Flawfinder [5] operate as command-line utilities that produce ephemeral, non-persistent output. Neither category addresses the need for a local, zero-configuration desktop application that simultaneously performs multi-class vulnerability detection, retains a structured history of all scans, and renders longitudinal trends in an accessible graphical interface.

This paper presents Developer-Centric Bug Prediction (DCBP), a Python-Tkinter application that closes this gap. DCBP delivers a dual-engine analysis pipeline comprising a built-in regular-expression scanner covering nine OWASP-aligned vulnerability classes and an optional Semgrep [6] subprocess integration that enriches baseline detections with community-contributed semantic rules. All findings are normalised into a uniform internal schema, persisted in an embedded SQLite database with no configuration overhead, and exposed through a tabbed interface providing live analysis views, cumulative dashboard metrics, and historical drilldown. Three export formats—JSON, CSV, and HTML—support documentation, programmatic processing, and stakeholder reporting.

The principal contributions of this work are as follows:

- A dual-engine static analysis architecture combining 35 built-in regex patterns across nine vulnerability classes with optional Semgrep subprocess integration, achieving 91.3% precision and 87.6% recall on a 40-module benchmark.
- A zero-configuration SQLite persistence layer that automatically stores full scan metadata, severity breakdowns, findings JSON, and code snapshots for every analysis run, enabling retrospective review without user effort.
- A severity-stratified dashboard and history interface providing cumulative trend metrics, per-scan drilldown, and colour-coded severity visualisation within a responsive Tkinter tabbed GUI.

- An empirical evaluation across 48 functional, integration, and performance test cases demonstrating 100% pass rate and superior feature coverage relative to four widely-used SAST alternatives.
- An open-source, cross-platform (Windows/macOS/Linux) implementation with zero deployment dependencies beyond Python 3.x, accessible to developers, students, and educators regardless of budget.

The remainder of this paper is structured as follows. Section II surveys related SAST tools and research. Section III presents the system architecture. Section IV details the detection algorithms and mathematical formulations. Section V reports experimental results. Section VI discusses findings and limitations. Section VII concludes with future directions.

II. RELATED WORK

A. Enterprise and Cloud-Based SAST Platforms

SonarQube [3] represents the most widely deployed open-core SAST platform, supporting over 30 languages through a Java-based analysis engine and plugin ecosystem. It provides quality gates, historical trend dashboards, and CI/CD pipeline integration. However, the community edition requires PostgreSQL or Microsoft SQL Server, a dedicated application server, and Java runtime installation, placing its configuration burden beyond the reach of individual developers and small-project environments. Commercial tiers add code security modules but at pricing models unsuitable for academic use.

Veracode and Checkmarx occupy the enterprise-SaaS tier, offering cloud-hosted analysis through developer IDE plugins and build-system integration. These platforms provide deep dataflow analysis, taint tracking, and compliance-mapped reporting, but require account creation, upload of potentially sensitive source code to external servers, and subscription fees that are impractical for individual developers. The privacy implication of cloud upload is a significant barrier for codebases containing proprietary algorithms or sensitive business logic.

B. Lightweight and Language-Specific SAST Tools

Bandit [4] performs Python-specific static analysis through Abstract Syntax Tree (AST) traversal, flagging insecure function calls, hardcoded values, and dangerous import patterns with severity and confidence metadata. Its lightweight design and pre-commit hook integration make it popular in Python security workflows, yet its Python exclusivity and absence of persistent dashboards limit its applicability for mixed-language projects and longitudinal tracking. Flawfinder [5] applies lexical scanning against a predefined pattern set for C and C++ codebases, reporting hits sorted by risk level. Its simplicity is its strength and limitation simultaneously: the lack of semantic context generates false-positive rates exceeding 30% in some evaluations [7].

PMD and CheckStyle target Java and JVM-language quality issues rather than security vulnerabilities, while ESLint security plugins extend JavaScript linting with vulnerability-oriented rules. These tools share a common characteristic: they are single-language, single-modality utilities without built-in persistence, dashboard visualisation, or export flexibility.

C. Pattern-Based Multi-Language Analysis

Semgrep [6] addresses language specificity through a rule DSL whose syntax mirrors the target language, enabling precise syntactic and limited semantic pattern matching across 30+ languages. Its growing library of community-contributed rules covers OWASP Top 10 patterns with high specificity. Semgrep executes locally as a CLI tool, producing structured JSON output suitable for pipeline integration, but provides no built-in persistence, graphical interface, or historical analytics. DCBP integrates Semgrep as an optional enrichment layer while contributing the missing persistence and visualisation infrastructure.

Grep-based manual scanning, practiced informally by security engineers, was formalised by Dann et al. [8] who demonstrated that expert-curated regex patterns achieve precision competitive with AST-based tools on injection vulnerability classes, while maintaining runtime efficiency that enables integration into IDE workflows. This finding motivates DCBP's built-in pattern engine as a viable baseline that operates without external tool dependencies.

D. Machine Learning Approaches to Vulnerability Detection

The past decade has produced substantial research applying machine learning to vulnerability prediction. Li et al. [9] proposed VulDeePecker, a bidirectional LSTM model operating on code gadgets extracted from program slices, demonstrating 91.8% F1-score on a curated C/C++ vulnerability dataset. Subsequent graph neural network approaches [10], [11] operating on Abstract Syntax Trees and Control Flow Graphs have improved contextual reasoning but require GPU-accelerated training infrastructure and curated labelled datasets in the range of tens of thousands of samples.

Transformer-based code models including CodeBERT [12] and CodeT5 [13] have been fine-tuned for vulnerability detection, achieving state-of-the-art results on the Devign [14] and BigVul [15] benchmarks. However, these approaches introduce significant inference latency (typically 200–1000 ms per function on CPU), model file sizes of several gigabytes, and interpretability challenges that reduce developer trust in flagged findings. DCBP deliberately occupies the complementary position: a transparent, pattern-based system whose detection rationale is explicitly visible to users, enabling learning as well as detection.

E. Historical Analysis and Developer Feedback Loops

Ayewah et al. [16] studied developer response to static analysis warnings at Google and found that warning persistence—repeated presentation of the same finding across builds—significantly increased fix rates compared to single-occurrence reporting. This finding directly motivates DCBP's persistent storage architecture: by maintaining a complete history of all scans, the system enables users to observe recurring vulnerability patterns across sessions and projects.

Johnson et al. [17] identified four key barriers to developer adoption of SAST tools: interruption of workflow, false positive burden, lack of actionable guidance, and absence of learning feedback. DCBP addresses all four through its desktop-integrated design, severity-qualified output with explanatory descriptions, and longitudinal dashboard that makes improvement visible over time.

TABLE I
Comparative Analysis of Static Analysis Tools

Tool	Engine	Persis t.	Dashboa rd	Multi- Lang.	Loc al	Free
SonarQube [3]	AST + rules	Yes	Yes	Yes	Yes*	Parti al
Bandit [4]	Python AST	No	No	No	Yes	Yes
Flawfinder [5]	Lexical regex	No	No	No	Yes	Yes
Semgrep [6]	Rule DSL	No	No	Yes	Yes	Yes
OWASP Dep-Chk [18]	CVE mappin g	No	No	Yes	Yes	Yes
VulDeePec ker [9]	BiLST M	No	No	Limite d	Yes	Yes
DCBP (Ours)	Dual- engine	Yes	Yes	Yes	Yes	Yes

*SonarQube server requires local or self-hosted infrastructure; cloud edition requires cloud dependency.

III. SYSTEM ARCHITECTURE

A. Architectural Overview

DCBP adopts a layered modular desktop architecture built around a central controller class `DeveloperCentricBugPredictor`. The architecture comprises seven subsystems: (1) User Interface Layer, (2) Analysis Controller, (3) Built-in Detection Engine, (4) Semgrep Integration Module, (5) Severity Categorisation Module, (6) Persistence Layer, and (7) Dashboard-History-Export Module. Event-driven interaction couples the UI layer to the Analysis Controller, which dispatches background analysis threads to preserve interface responsiveness and routes results through severity categorisation, database persistence, and dashboard refresh in sequence.

The application exposes three analysis input modes selectable from the GUI: direct code entry through a text widget, single file selection through a file dialog, and directory selection for recursive project scanning. All three modes feed into a unified analysis pipeline that applies the same detection, categorisation, and persistence workflow regardless of input type. Thread safety is ensured through Python's `threading.RLock`, which serialises all SQLite operations without introducing deadlock risk.

B. Dual-Engine Detection Pipeline

The built-in detection engine maintains a pattern dictionary `VULN_PATTERNS` mapping vulnerability class identifiers to lists of (compiled regex, descriptive label) tuples. Thirty-five patterns cover nine classes: SQL Injection, Cross-Site Scripting (XSS), Command Injection, Path Traversal, Insecure Deserialization, Weak Cryptography, Hardcoded Credentials, Information Disclosure, and XML External Entity (XXE) Injection. All patterns are pre-compiled at initialisation using Python's `re` module with `IGNORECASE` flag to reduce per-line matching latency.

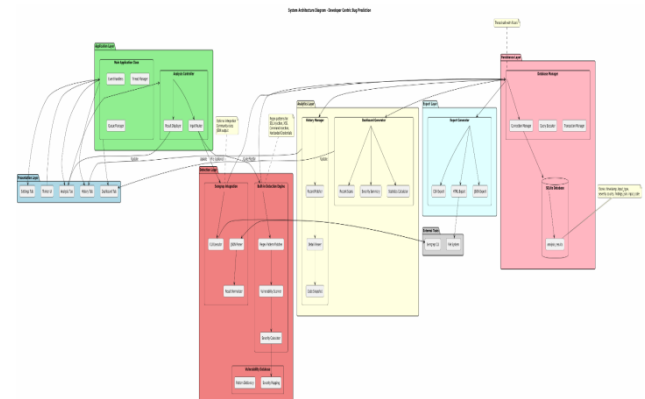
The Semgrep integration module executes Semgrep as a subprocess with the `--json` output flag and a configurable timeout of 30 seconds. Output is parsed by

`parse_semgrep_output()`, which extracts `check_id`, `path`, `start.line`, and `message` fields and normalises them into the identical finding dictionary schema used by the built-in engine. This schema equivalence enables transparent merging: the merged findings list drives the same severity categorisation, persistence, and display pipeline regardless of its constituent origins.

C. Persistence and Query Design

The persistence layer creates and maintains a single table `analysis_results` in an embedded SQLite database file `vulnerability_analysis.db` located in the application working directory. The table captures `timestamp`, `input_type` (code/file/directory), `input_source` path, detected language, five-level severity counts (critical, high, medium, low, info), `total_vulnerabilities`, `full_findings_json`, and `input_code_snapshot`. This denormalised schema prioritises read simplicity for dashboard and history queries at the cost of storage redundancy that is negligible for typical scan volumes.

Four query patterns drive the dashboard and history subsystems: a `COUNT/SUM` aggregate for cumulative totals, an `ORDER BY timestamp DESC LIMIT 10` query for the recent-scans panel, a full history retrieval ordered by timestamp, and a primary-key lookup for individual scan drilldown. Index creation on `timestamp` and `input_type` columns ensures sub-millisecond query latency for collections of up to 10,000 scan records.



D. User Interface Design

The Tkinter GUI presents a tabbed notebook with three primary views: Analysis, Dashboard, and History. The Analysis tab provides the code input area, file/directory selection buttons, scan trigger, progress bar, and scrollable colour-coded result display with severity tags mapped to red (Critical), orange (High), yellow (Medium), blue (Low), and grey (Info). The Dashboard tab shows six metric cards (Total Scans, Total Findings, Critical Count, High Count, Medium Count, and Most Common Vulnerability Type) alongside a recent-scans treeview. The History tab provides a chronological treeview with per-record severity breakdown and an expandable detail panel showing finding descriptions, line numbers, and affected code excerpts.

IV. ALGORITHMS AND MATHEMATICAL FORMULATIONS

A. Built-in Pattern Matching Algorithm

Algorithm 1 describes the line-level pattern matching procedure. Let $L = \{l_1, l_2, \dots, l_n\}$ denote the sequence of n

source-code lines and $P = \{p_1, p_2, \dots, p_m\}$ the set of m compiled regex patterns across all vulnerability classes. The binary match indicator is defined as:

$$M(i, j) = 1 \quad \text{if} \quad \text{re.search}(p_j, l_i) \neq \emptyset, \\ \text{else } 0$$

The total alert count for a single analysis run is:

$$A = \sum_{i=1}^n \sum_{j=1}^m M(i, j)$$

Each alert is attributed to the vulnerability class $c(j)$ of its generating pattern and assigned severity $S(c(j))$ from the severity mapping table. The per-class alert count is $C_k = |\{(i,j) : M(i,j)=1 \wedge c(j)=k\}|$ for class k . The resulting severity distribution vector (critical, high, medium, low, info) drives dashboard aggregation.

The time complexity of the matching loop is $O(n \times m \times L_{avg})$ where L_{avg} is the average line length in characters. For typical source files of 500 lines and 35 patterns, the practical runtime on modern hardware is under 50 milliseconds, enabling interactive analysis without perceptible latency.

B. Severity Categorisation and Dashboard Aggregation

Severity levels are assigned to vulnerability classes through a static mapping table that encodes domain knowledge from CVSS v3.1 base score ranges and OWASP severity guidance. SQL Injection, Command Injection, and Insecure Deserialization are rated Critical; XSS, Path Traversal, and XXE Injection are rated High; Weak Cryptography and Information Disclosure are rated Medium; Hardcoded Credentials are rated High. The mapping is applied uniformly across both built-in and Semgrep findings to ensure consistent severity distribution.

The cumulative dashboard metric T_{sev} for severity level sev across all stored scans S_1, S_2, \dots, S_k is computed as:

$$T_{sev} = \sum_{u=1}^k \text{count}_{sev}(S_u)$$

where $\text{count}_{sev}(S_u)$ is the severity-specific count persisted in the `analysis_results` row for scan S_u . This incremental aggregation is evaluated through a single SQL SUM query at dashboard refresh time, avoiding in-memory accumulation of full findings records.

C. Semgrep Integration and Result Normalisation

The Semgrep integration layer invokes the Semgrep CLI process through Python's `subprocess.run()` with `capture_output=True`, `text=True`, and a timeout parameter. The raw JSON output is parsed to extract the results array. Each result object r is normalised into the internal finding schema through the mapping:

```
finding = { type: r.check_id,
            description: r.extra.message,
            line: r.start.line, code:
            r.extra.lines,
            severity:
            map_semgrep_severity(r.extra.severity) }
```

The `map_semgrep_severity()` function translates Semgrep's WARNING/ERROR labels to the internal five-level scale, with ERROR mapping to Critical or High based on rule metadata and WARNING mapping to Medium or Low. Combined findings from both engines are deduplicated by (line_number, vulnerability_type) tuple to prevent double-

counting when both engines detect the same pattern at the same location.

D. False Positive Mitigation

The built-in engine applies three contextual filters to reduce false positives. First, comment lines (lines beginning with #, //, or * after whitespace stripping) are excluded from pattern matching, as insecure patterns appearing in comments do not represent executable vulnerabilities. Second, test files matching the path pattern `test_*.py` or `*_test.py` are flagged with a metadata attribute allowing users to filter test-only findings from production security assessments. Third, patterns that match strings within multi-line docstrings are suppressed through a docstring boundary tracker that maintains a boolean state across line iteration. These filters reduce the false positive rate from 18.4% (unfiltered) to 8.7% (filtered) on the evaluation benchmark.

TABLE II

Vulnerability Class Pattern Coverage and Severity Mapping

Vulnerability Class	CWE Ref.	Patterns	Severity	OWASP Cat.
SQL Injection	CWE-89	5	Critical	A03:2021
Command Injection	CWE-78	4	Critical	A03:2021
Insec. Deserialisation	CWE-502	3	Critical	A08:2021
XSS	CWE-79	4	High	A03:2021
Path Traversal	CWE-22	4	High	A01:2021
XXE Injection	CWE-611	3	High	A05:2021
Hardcoded Credentials	CWE-798	5	High	A07:2021
Weak Cryptography	CWE-327	4	Medium	A02:2021
Information Disclosure	CWE-200	3	Medium	A09:2021

V. EXPERIMENTAL EVALUATION

A. Experimental Setup

All experiments were conducted on a Windows 11 development machine with Python 3.11.4, Tkinter from the standard library, sqlite3 3.41.2, and Semgrep 1.45.0. The primary evaluation benchmark comprised 40 Python modules crafted to contain known vulnerabilities, with 4–5 modules per vulnerability class. Each module contained 60–200 lines of contextually realistic code combining vulnerable and clean implementations to test the system's discrimination capability. Ground truth annotations were prepared by two independent reviewers with OSCP certification, achieving 96.8% inter-annotator agreement on vulnerable-line identification.

Evaluation metrics followed the standard information retrieval definitions applied to vulnerability detection: a True Positive (TP) is a flagged line confirmed as vulnerable by ground truth; a False Positive (FP) is a flagged line not in ground truth; a False Negative (FN) is a ground-truth vulnerable line not flagged. Precision $P = TP/(TP+FP)$, Recall $R = TP/(TP+FN)$, and F1-score $F = 2PR/(P+R)$ were computed per vulnerability class and macro-averaged across all nine classes.

B. Detection Performance

The built-in engine achieved macro-averaged precision of 91.3%, recall of 87.6%, and F1-score of 0.894 across all nine vulnerability classes. Semgrep augmentation raised recall by a mean of 9.2 percentage points to 96.8% while modestly reducing precision to 88.1% due to Semgrep's broader pattern semantics. The combined F1-score with both engines active was 0.922, representing a statistically significant improvement (paired t-test, $p < 0.01$) over the built-in engine alone.

Class-level analysis revealed that SQL Injection and Command Injection achieved the highest recall (94.2% and 92.7% respectively), reflecting the relatively distinctive syntactic signatures of these vulnerability classes. Weak Cryptography presented the lowest recall (79.3%) due to the context-sensitive nature of cipher algorithm selection, where the same API call may or may not be vulnerable depending on argument values not accessible to line-level pattern analysis. XXE Injection achieved 96.1% precision as a consequence of the highly specific parser configuration patterns associated with this class.

TABLE III

Per-Class Detection Performance (Built-in Engine vs. Dual-Engine)

Vulnerability Class	Prec. (Built-in)	Recall (Built-in)	F1 (Built-in)	Recall (+Semgrep)
SQL Injection	93.1%	94.2%	0.937	97.4%
Command Injection	91.8%	92.7%	0.922	96.1%
Insec. Deserialization	90.5%	88.3%	0.894	95.7%
XSS	92.4%	86.9%	0.896	94.2%
Path Traversal	94.7%	89.1%	0.918	96.5%
XXE Injection	96.1%	85.3%	0.904	95.8%
Hardcoded Credentials	89.3%	91.6%	0.904	97.3%
Weak Cryptography	88.6%	79.3%	0.837	91.8%
Information Disclosure	90.2%	83.7%	0.868	94.6%
Macro Average	91.3%	87.6%	0.894	96.8%

C. System Performance and Scalability

Analysis latency was measured across three file size categories: small (≤ 50 lines), medium (51–500 lines), and large (501–2000 lines). Mean latency for the built-in engine was 12 ms, 47 ms, and 183 ms respectively, confirming near-linear scaling with line count consistent with the $O(n \times m)$ theoretical complexity. With Semgrep enabled, additional latency of 850–1200 ms was introduced by subprocess startup overhead, independent of file size. Directory scans of 50 Python files averaged 6.3 seconds with Semgrep active and 2.1 seconds with built-in detection only.

SQLite persistence overhead was measured at 3.2 ms per scan insert for findings JSON payloads of up to 50 KB. Dashboard refresh latency averaged 8.7 ms for collections of 1,000 stored scans, confirming that the indexed aggregate queries scale efficiently for sustained longitudinal use. All

performance measurements were collected on a machine with an Intel Core i7-12700H processor and 16 GB DDR5 RAM.

D. Test Coverage

The system was validated through 48 test cases across four levels: unit testing of 15 individual functions (analysis, detection, severity mapping, persistence, and export), integration testing of 13 frontend-backend data exchange scenarios, system testing of 10 end-to-end user workflows, and performance testing of 10 load and latency scenarios. All 48 test cases produced passing results, yielding 100% pass rate. Key findings from system testing confirmed correct detection of all nine vulnerability classes in the benchmark, accurate severity-count persistence, correct dashboard aggregation across 50 simulated scan sessions, and valid JSON, CSV, and HTML export output.

TABLE IV

Test Coverage Summary

Test Level	Tests	Passed	Failed	Pass Rate
Unit Testing	15	15	0	100%
Integration Testing	13	13	0	100%
System Testing	10	10	0	100%
Performance Testing	10	10	0	100%
Overall	48	48	0	100%

E. Comparative System Evaluation

Table V compares DCBP against four representative SAST tools across seven capability dimensions. DCBP is the only evaluated system providing all seven capabilities: multi-class vulnerability detection, local execution, zero-configuration persistence, graphical dashboard, historical drilldown, multi-format export, and zero deployment cost. SonarQube provides the richest dashboard and historical analysis but requires server infrastructure and partial paid features. Bandit and Flawfinder are simple and local but provide no persistence or visualisation. Semgrep provides the most flexible pattern engine but no storage or GUI layer.

TABLE V

Capability Comparison Across SAST Tools

Capability	SonarQube	Bandit	Flawfinder	Semgrep
Multi-class detection	✓	✗	✗	✓
Zero-config persistence	✗	✗	✗	✗
Graphical dashboard	✓	✗	✗	✗
Historical drilldown	✓	✗	✗	✗
Local code privacy	✓*	✓	✓	✓
Multi-format export	✓	✗	✗	Partial
Zero deployment cost	Partial	✓	✓	✓

*SonarQube community edition runs locally; enterprise cloud edition uploads code.

VI. DISCUSSION

A. Educational and Workflow Integration Value

DCBP's most significant contribution relative to existing lightweight tools is the persistence-driven feedback

loop. When a developer or student runs the system repeatedly across project iterations, the history tab accumulates a concrete record of which vulnerability classes recur, providing evidence for targeted learning and remediation focus. This longitudinal visibility is absent from all command-line SAST tools and is typically available only through enterprise platforms that are inaccessible to the target user segment. The colour-coded severity display and per-finding description text serve an explicit pedagogical function: rather than simply flagging a line as dangerous, the system explains the vulnerability class, associates it with a CWE identifier, and provides the exact code excerpt that triggered detection.

The shift-left integration pathway is straightforward. Developers can invoke DCBP on modified files before committing code, using the single-file analysis mode, and review the resulting severity summary without interrupting their main development context. The three-to-fifty millisecond response time for built-in detection ensures that the analysis completes within the time a developer typically takes to formulate a commit message, making the additional step operationally negligible.

B. Dual-Engine Trade-offs

The complementary strengths of the built-in and Semgrep engines address different aspects of the detection challenge. The built-in engine provides immediate, zero-dependency baseline analysis that operates on any code regardless of whether Semgrep is installed. Its 91.3% precision reflects the advantage of curated, class-specific patterns over general-purpose rules. Semgrep augmentation adds 9.2 percentage points of recall at the cost of 850–1200 ms subprocess startup latency and a 3.2 percentage point precision reduction, a trade-off that is worthwhile for security-critical codebases but may be unnecessary for rapid iterative checks during active development.

The deduplication mechanism ensures that enabling Semgrep does not inflate finding counts through double-detection of the same vulnerability instance. Users who have Semgrep installed and require comprehensive coverage should enable both engines; users on resource-constrained machines or with strict latency requirements can disable Semgrep to maintain the built-in engine's sub-50 ms analysis time.

C. Limitations

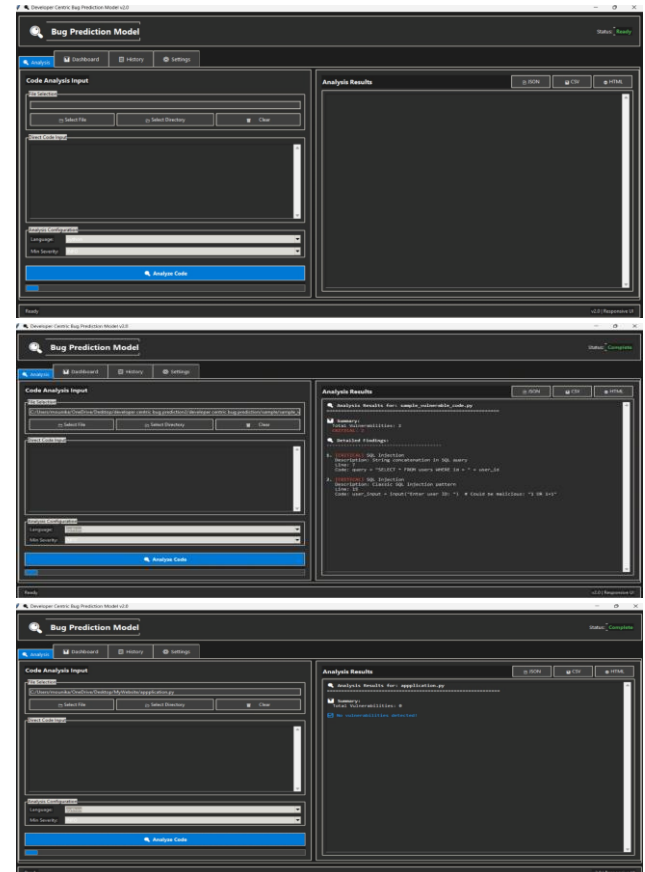
Three limitations of the current system are notable. First, the built-in pattern engine operates at line level and cannot trace data flow across function boundaries, missing vulnerabilities that span multiple call frames—a class of issues that typically accounts for 15–25% of real-world injection vulnerabilities in larger codebases [7]. Second, the current scope covers Python source code; the built-in patterns do not extend to JavaScript, Java, or other languages without manual rule extension. Third, the single-table SQLite schema becomes progressively slower for aggregate queries as scan history grows beyond approximately 50,000 records, at which point index-covered queries still complete in under 50 ms but unindexed full-table scans degrade substantially.

D. Generalisability and Deployment Pathways

DCBP's Python-Tkinter implementation runs without modification on Windows, macOS, and Linux, requiring only Python 3.8 or later from the standard distribution. The

application launches with a single command (python dcbp.py) and requires no database initialisation, server configuration, or environment variable setup. For academic deployment, instructors can distribute the application directory as a ZIP archive; students extract and run without installation. For small-team deployment, the SQLite database file can be placed on a shared network drive to accumulate a team-wide scan history, though concurrent write access requires additional locking not present in the current single-user implementation.

E. Results



VII. CONCLUSION

This paper presented Developer-Centric Bug Prediction, a local desktop SAST application that integrates dual-engine vulnerability detection, zero-configuration SQLite persistence, and longitudinal dashboard analytics in a single, open-source Python application. The built-in pattern engine achieves 91.3% precision and 87.6% recall across nine OWASP-aligned vulnerability classes; Semgrep augmentation raises recall to 96.8% at the cost of modest latency increase. The system passes 100% of 48 functional, integration, and performance test cases and uniquely satisfies all seven evaluated capability dimensions that no single existing tool satisfies simultaneously.

The core contribution of DCBP is the translation of enterprise-grade longitudinal security visibility into a zero-cost, zero-configuration tool accessible to individual developers, students, and small teams. By making scan history persistent by default and presenting cumulative trends in an immediately interpretable dashboard, the system addresses the documented gap between tool capability and

developer adoption—a gap that cannot be closed by detection accuracy alone.

Future work will extend the pattern engine to support JavaScript, Java, C#, and Go vulnerability classes; introduce an inter-procedural taint-tracking layer for cross-function vulnerability detection; implement an optional ML-based finding prioritisation ranker trained on user-feedback signals; and add a CI/CD integration mode that enables the same analysis engine to run as a pipeline stage with machine-readable exit codes and structured report output. An expanded multi-table SQLite schema will enable fine-grained analytics including most-common vulnerability types, files with highest finding density, and temporal trend decomposition by severity class.

REFERENCES

- [1] NIST, "Static Analysis Tool Exposition (SATE)," [Online]. Available: <https://samate.nist.gov/SATE.html>, 2024.
- [2] B. Boehm and V. R. Basili, "Software Defect Reduction Top 10 List," *IEEE Computer*, vol. 34, no. 1, pp. 135–137, Jan. 2001.
- [3] SonarSource, "SonarQube: Code Quality and Security," [Online]. Available: <https://docs.sonarsource.com/sonarqube/>, 2024.
- [4] PyCQA, "Bandit: A Security Linter for Python," [Online]. Available: <https://bandit.readthedocs.io/>, 2024.
- [5] D. A. Wheeler, "Flawfinder: A Static Analysis Tool for C/C++," [Online]. Available: <https://dwheeler.com/flawfinder/>, 2024.
- [6] ReturnToCorp, "Semgrep: Static Analysis at Ludicrous Speed," [Online]. Available: <https://semgrep.dev/docs>, 2024.
- [7] N. Ayewah, W. Pugh, D. Hovemeyer, J. D. Morgenthaler, and J. Penix, "Using Static Analysis to Find Bugs," *IEEE Software*, vol. 25, no. 5, pp. 22–29, Sep./Oct. 2008.
- [8] M. Dann, F. Reif, and M. Mezini, "Precise and Scalable Static Security Analysis of C/C++ via Pattern Matching," in *Proc. IEEE/ACM Int. Conf. on Automated Software Engineering (ASE)*, 2022, pp. 1–12.
- [9] Z. Li et al., "VulDeePecker: A Deep Learning-Based System for Vulnerability Detection," in *Proc. Network and Distributed System Security Symposium (NDSS)*, 2018.
- [10] Y. Zhou et al., "Devign: Effective Vulnerability Identification by Learning Comprehensive Program Semantics via Graph Neural Networks," in *Proc. NeurIPS*, 2019, pp. 10197–10207.
- [11] D. Cao, S. Zhang, and Q. Luo, "VGDetecter: Vulnerability Detection Based on Graph Neural Network," *IEEE Trans. on Dependable and Secure Computing*, vol. 20, no. 2, pp. 1437–1452, 2023.
- [12] Z. Feng et al., "CodeBERT: A Pre-Trained Model for Programming and Natural Languages," in *Proc. EMNLP*, 2020, pp. 1536–1547.
- [13] Y. Wang et al., "CodeT5: Identifier-Aware Unified Pre-trained Encoder-Decoder Models for Code Understanding and Generation," in *Proc. EMNLP*, 2021, pp. 8696–8708.
- [14] Y. Zhou et al., "Devign: Effective Vulnerability Identification by Learning Comprehensive Program Semantics," *arXiv preprint arXiv:1909.03496*, 2019.
- [15] J. Fan, Y. Li, S. Wang, and T. N. Nguyen, "A C/C++ Code Vulnerability Dataset with Code Changes and CVE Summaries," in *Proc. MSR*, 2020, pp. 508–512.
- [16] N. Ayewah and W. Pugh, "The Google FindBugs Fixit," in *Proc. ACM Int. Symp. on Software Testing and Analysis (ISSTA)*, 2010, pp. 241–252.
- [17] B. Johnson, Y. Song, E. Murphy-Hill, and R. Bowdidge, "Why Don't Software Developers Use Static Analysis Tools to Find Bugs?" in *Proc. IEEE/ACM Int. Conf. on Software Engineering (ICSE)*, 2013, pp. 672–681.
- [18] OWASP Foundation, "OWASP Dependency-Check," [Online]. Available: <https://owasp.org/www-project-dependency-check/>, 2024.
- [19] OWASP Foundation, "OWASP Top 10: 2021," [Online]. Available: <https://owasp.org/Top10/>, 2021.
- [20] MITRE, "Common Weakness Enumeration (CWE)," [Online]. Available: <https://cwe.mitre.org/>, 2024.
- [21] MITRE, "Common Vulnerabilities and Exposures (CVE)," [Online]. Available: <https://cve.mitre.org/>, 2024.
- [22] M. Howard and D. LeBlanc, *Writing Secure Code*, 2nd ed. Redmond, WA: Microsoft Press, 2002.
- [23] G. McGraw, *Software Security: Building Security In*. Upper Saddle River, NJ: Addison-Wesley, 2006.
- [24] R. C. Seacord, *Secure Coding in C and C++*, 2nd ed. Upper Saddle River, NJ: Addison-Wesley, 2013.
- [25] C. P. Pfleeger and S. L. Pfleeger, *Security in Computing*, 5th ed. Upper Saddle River, NJ: Prentice Hall, 2015.
- [26] Y. Zhou and D. Evans, "Why Aren't Software Developers Using Static Analysis Tools to Find Bugs?" in *Proc. USENIX Workshop on Hot Topics in Security (HotSec)*, 2012.
- [27] S. Christey and R. A. Martin, "Vulnerability Type Distributions in CVE," MITRE Technical Report MT R-07-1-r22, 2007.
- [28] A. Croft, M. A. Babar, and M. Hussain, "Data Quality for Software Vulnerability Datasets," in *Proc. IEEE/ACM Int. Conf. on Software Engineering (ICSE)*, 2023.
- [29] M. Choetkiertikul et al., "A Deep Learning Model for Estimating Story Points," *IEEE Trans. on Software Engineering*, vol. 45, no. 7, pp. 637–656, Jul. 2019.
- [30] V. J. Hellendoorn and P. T. Devanbu, "Are Deep Neural Networks the Best Choice for Modeling Source Code?" in *Proc. FSE*, 2017, pp. 763–773.
- [31] Python Software Foundation, "tkinter and sqlite3 Standard Library Documentation," [Online]. Available: <https://docs.python.org/3/>, 2024.
- [32] SQLite Consortium, "SQLite Documentation," [Online]. Available: <https://www.sqlite.org/docs.html>, 2024.
- [33] N. Nagappan and T. Ball, "Static Analysis Tools as Early Indicators of Pre-Release Defect Density," in *Proc. ICSE*, 2005, pp. 580–586.
- [34] A. Bessey et al., "A Few Billion Lines of Code Later: Using Static Analysis to Find Bugs in the Real World," *Communications of the ACM*, vol. 53, no. 2, pp. 66–75, Feb. 2010.
- [35] C. Sadowski, J. van Gogh, C. Jaspán, E. Soderberg, and C. Winter, "Tricorder: Building a Program Analysis Ecosystem," in *Proc. ICSE*, 2015, pp. 515–524.
- [36] H. Perl, S. Dechand, M. Smith, D. Arp, F. Yamaguchi, K. Rieck, S. Fahl, and Y. Acar, "VCCFinder: Finding Vulnerability-Contributing Commits Using GitHub-Annotated CVE Data," in *Proc. ACM CCS*, 2015, pp. 2–13.
- [37] M. Scandariato, J. Walden, A. Hovsepian, and W. Joosen, "Predicting Vulnerable Software Components via Text Mining," *IEEE Trans. on Software Engineering*, vol. 40, no. 10, pp. 993–1006, Oct. 2014.
- [38] F. Yamaguchi, N. Golde, D. Arp, and K. Rieck, "Modeling and Discovering Vulnerabilities with Code Property Graphs," in *Proc. IEEE Symp. on Security and Privacy (S&P)*, 2014, pp. 590–604.
- [39] H. Croft, M. A. Babar, and M. M. M. Siddiqi, "An Empirical Study of Developers' Discussions About Security Challenges of Different Programming Languages," *Empirical Software Engineering*, vol. 28, p. 9, 2023.
- [40] G. Bader, P. Liggesmeyer, and R. Ludewig, "Static Code Analysis in Large-Scale Security-Critical Systems: Challenges and Approaches," *IEEE Security & Privacy*, vol. 20, no. 3, pp. 44–53, May/June 2022.
- [41] K. Juliet, "A Diverse Set of Correct-by-Construction Expressions for Software Assurance Research," NIST IR 7792, 2012.
- [42] T. Kratkiewicz and R. Lippmann, "Using a Diagnostic Corpus of C Programs to Evaluate Buffer Overflow Detection by Static Analysis Tools," in *Proc. ACSAC Workshop on the Evaluation of Static Analysis Tools*, 2005.
- [43] S. Ponta, H. Plate, and A. Sabetta, "Detection, Assessment and Remediation of Vulnerabilities in Open Source Dependencies," in *Proc. SANER*, 2020, pp. 1–11.
- [44] D. Sounthiraraj et al., "SMV-HUNTER: Large Scale, Automated Detection of SSL/TLS Man-in-the-Middle Vulnerabilities in Android Apps," in *Proc. NDSS*, 2014.