



Human-in-the-Loop Secure Code Synthesis: Integrating Security Heuristics in AI Code Generation

Mohan Siva Krishna Konakanchi
Independent Researcher.

Received On: 10/08/2025

Revised On: 22/08/2025

Accepted On: 15/09/2025

Published On: 03/10/2025

Abstract - The proliferation of Large Language Models (LLMs) as code generation assistants has revolutionized software development but simultaneously introduced a new vector for security vulnerabilities. These models, trained on vast repositories of public code, often replicate insecure patterns, leading to the generation of code susceptible to common exploits. To address this challenge, we propose a novel Human-in-the-Loop (HITL) framework for secure code synthesis that synergistically combines real-time static analysis with LLM-based code generation directly within the Integrated Development Environment (IDE). Our system, named *SecurifyAI*, employs an iterative refinement loop where code snippets generated by an LLM are immediately scrutinized by a lightweight, high-speed Static Application Security Testing (SAST) engine. The identified potential vulnerabilities are then translated into contextual feedback to guide the LLM in regenerating a more secure version of the code. This proactive approach shifts security from a post-development afterthought to an integral part of the code creation process. Furthermore, we introduce a trust metric-based federated learning (FL) framework to continuously improve the underlying LLM's security awareness across distributed, privacy-sensitive environments. This FL approach ensures integrity and accountability by weighting contributions from different clients based on a calculated trust score. Finally, we propose a formal model to quantify and optimize the inherent trade-off between the explainability of security warnings and the performance of the code generation system. Our experimental results, conducted on a curated dataset of insecure code patterns, demonstrate that *SecurifyAI* reduces the incidence of common vulnerabilities (CWEs) by over 80% compared to baseline LLM assistants, while maintaining acceptable latency and improving developer code acceptance rates through clear, actionable feedback.

Keywords - Secure Code Generation, Large Language Models, Human-in-the-Loop, Static Analysis, Federated Learning, Explainable AI, Application Security.

1. Introduction

The advent of Large Language Models (LLMs) such as OpenAI's Codex [1] and Google's AlphaCode has

profoundly impacted the software development lifecycle. Integrated into developer environments, these AI-powered assistants accelerate coding tasks, suggest boilerplate code, and even author complex algorithms from natural language prompts. This paradigm shift, however, carries significant security implications. LLMs are trained on massive datasets of publicly available source code, including codebases with latent or explicit security flaws. Consequently, these models often inadvertently learn and propagate insecure coding practices, generating snippets. This work was not funded by any organization. susceptible to vulnerabilities like SQL injection, Cross-Site Scripting (XSS), and buffer overflows [2], [3]. The conventional approach to software security, relying on post-hoc analysis through Static Application Security Testing (SAST) and Dynamic Application Security Testing (DAST) tools, is reactive. By the time these tools identify vulnerabilities, the insecure code is already integrated into the codebase, making remediation more costly and time-consuming. There is a pressing need for a proactive, "secure-by-default" paradigm that integrates security directly into the moment of code creation.

In this paper, we present *SecurifyAI*, a novel Human-in-the-Loop (HITL) framework designed to synthesize secure code by default. Our approach creates a tight, real-time feedback loop between an LLM code generator and a lightweight SAST engine within the IDE. When a developer requests code, the LLM generates an initial version, which is immediately analyzed for security weaknesses. The SAST findings are not merely presented as warnings but are programmatically transformed into a refined prompt that instructs the LLM to correct its output. The developer is presented with the secured code and a clear explanation of the mitigated risk, empowering them to make informed decisions (the "human in the loop"). To ensure that the underlying LLM continuously learns from these interactions without compromising user privacy or proprietary code, we propose a trust-metric based federated learning (FL) framework. Organizations can deploy local instances of *SecurifyAI*, and the security-centric model improvements (e.g., gradients from correcting vulnerabilities) are aggregated centrally. Our novel trust metric ensures the integrity of the global model by weighting contributions from each client based on the

quality and consistency of their provided feedback, thus mitigating risks of model poisoning. Finally, we recognize that developer adoption hinges on usability. An overly aggressive or opaque security tool can disrupt workflows. We address this by introducing a quantitative framework for managing the trade-off between the system's performance (latency, accuracy) and the explainability of its security interventions. This allows for tuning the system to provide the optimal level of detail in its explanations without overwhelming the developer or slowing down the coding process.

The primary contributions of this work are:

- A novel HITL framework, *SecurifyAI*, that integrates a real-time SAST feedback loop with an LLM to proactively generate secure code snippets within an IDE.
- The design of a trust metric-based federated learning architecture that enables continuous, privacy-preserving improvement of the secure coding model across decentralized clients.
- A formal methodology to quantify and optimize the trade-off between the explainability of security feedback and the overall performance of the code synthesis system.
- An empirical evaluation demonstrating that our approach significantly reduces the generation of vulnerable code compared to state-of-the-art AI coding assistants.

The remainder of this paper is structured as follows. Section II reviews related work in AI code generation and its security. Section III details the methodology behind *SecurifyAI*. Section IV describes our experimental setup and presents the results. Section V discusses the implications and limitations of our findings. Finally, Section VI concludes the paper and outlines directions for future research.

2. Related Work

Our research builds upon several distinct but converging fields: AI for code generation, the security of machine-generated code, Human-in-the-Loop systems, and federated learning for security.

2.1. AI for Code Generation

The use of deep learning for code generation has seen remarkable progress. Early models based on Recurrent Neural Networks (RNNs) and LSTMs showed promise in generating small code snippets [4]. The advent of the Transformer architecture [5] led to a significant leap in capability, culminating in large-scale models like GPT-3 and its derivatives, such as OpenAI's Codex [1], which powers GitHub Copilot. These models can translate natural language into complex code across various programming languages. While their productivity benefits are undisputed, their training on unfiltered web-

scale data makes them prone to security errors [6].

2.2. Security of AI-Generated Code

Several studies have highlighted the security risks associated with LLM-based code assistants. Pearce et al. [2] conducted a user study where participants using an AI assistant were more likely to produce insecure code than those without. Siddiq et al. [7] analyzed thousands of code suggestions from Copilot and found a significant percentage contained vulnerabilities from the MITRE CWE Top 25. These findings underscore the "insecure-by-default" nature of current models and motivate the need for integrated security guardrails. Existing solutions often propose post-generation security scanners, which is a reactive measure that we aim to improve upon with our proactive synthesis approach.

2.3. Human-in-the-Loop (HITL) Systems

HITL is a paradigm where a computational model and a human user interact to solve a problem, leveraging the strengths of both. It is widely used in data labeling, content moderation, and complex decision-making systems [8]. In the context of software engineering, HITL has been applied to debugging and program repair [9]. Our work applies the HITL concept to secure code generation, where the AI provides the initial solution, the SAST engine provides automated security expertise, and the human developer provides the final validation and contextual feedback, creating a virtuous cycle of improvement.

2.4. Federated Learning for Security

Federated Learning (FL) is a distributed machine learning technique that enables model training on decentralized data without data centralization [10]. It is particularly well-suited for privacy-sensitive domains like healthcare and finance. In cybersecurity, FL has been used for network intrusion detection and malware analysis [11]. We are the first, to our knowledge, to propose its application for building a globally robust secure code generation model. Our primary innovation in this area is the introduction of a trust metric to enhance the integrity and accountability of the federated learning process, addressing concerns about malicious clients providing poisoned model updates [12].

3. Methodology: The SecurifyAI Framework

The *SecurifyAI* framework is designed as a proactive, developer-centric system to synthesize secure code. Its architecture revolves around three core components: the Secure Code Synthesis Loop, the Trust Metric-based Federated Learning mechanism, and the Explainability-Performance Optimization model.

3.1. System Architecture and Secure Code Synthesis Loop

The workflow of *SecurifyAI* is embedded within the developer's IDE and operates in real-time. The process, illustrated as a sequence of steps, is as follows:

- Prompt Ingestion: The developer writes a natural lan-

guage comment or code stub (e.g., ‘// Function to validate user input and query the database’).

- **Initial Code Generation:** This prompt is sent to a fine-tuned LLM for code (e.g., a secured version of GPT-4 or a similar model), which generates a preliminary code snippet.
- **Real-time SAST Analysis:** The generated snippet is immediately passed to an integrated, lightweight SAST engine. This engine is optimized for speed and focuses on high-confidence, high-impact vulnerability patterns (e.g., OWASP Top 10) to minimize latency. It scans for issues like unsanitized database queries, improper handling of user input, or use of deprecated cryptographic functions.
- **Iterative Refinement via Feedback Prompting:** If the SAST engine detects a potential vulnerability, the system does not simply flag it. Instead, it programmatically constructs a new prompt that includes the original request, the flawed code, and a directive to fix the specific security issue. For example: “The user wants a function to query the database. The previous code ‘String Query’ average of the client model parameters W_i , with the weights being the normalized trust scores:

$$w_G^{(t+1)} = \frac{\sum_{j=1}^N T_j(t) \cdot w_j^{(t+1)}}{\sum_{j=1}^N T_j(t)}$$

= ”SELECT * FROM users WHERE name = ” + name

+ ””; ‘ is vulnerable to SQL Injection. Regenerate the code using parameterized queries or prepared statements to prevent this.”

- **Secure Code Presentation:** The LLM processes the refined prompt and generates a corrected version of the code. This secure snippet is then presented to the developer.
- **Human-in-the-Loop Feedback:** Alongside the code, an explanation is provided (e.g., “Used a PreparedStatement to prevent SQL Injection.”). The developer can accept the code, reject it, or manually edit it. This interaction (acceptance, rejection, edit distance) serves as a valuable feedback signal for future model training.

This entire loop is designed to complete in sub-second time to avoid disrupting the developer’s flow.

3.2. Trust Metric-based Federated Learning

To enable *SecurifyAI* to learn from a diverse range of coding contexts without centralizing proprietary source code, we employ a federated learning approach. Multiple clients (e.g., different organizations) run local instances of the system. The model improvements derived from the HITL feedback are shared, not the code itself. To ensure the integrity of the global model, we introduce a client-side Trust Metric, T_i , for each

client i .

The trust metric for a client i at a given training round t is defined as:

This mechanism ensures that clients providing high-quality, consistent, and effective security feedback have a greater influence on the global model, enhancing its robustness and security.

3.3. Quantifying the Explainability-Performance Trade-off

A crucial aspect of an HITL system is its usability. The explanations provided to the developer must be helpful but not overwhelming. We model the trade-off between explainability (X) and performance (P).

Performance (P) is a composite metric defined as:

$$P = w_1(1 - VGR) + w_2(CAR) - w_3(\text{Latency})$$

Where:

- VGR is the Vulnerability Generation Rate (fraction of generated snippets that are insecure).
- CAR is the Code Acceptance Rate by the developer.
- Latency is the end-to-end time for code generation.
- w_1, w_2, w_3 are importance weights.

Explainability (X) is quantified based on the utility of the security feedback. We measure this through user studies using a combination of the System Usability Scale (SUS) [13] for perceived usability and a task-based metric where we measure the time it takes for a developer to understand and correctly act upon a security warning. A higher X value corresponds to clearer, more actionable explanations.

The trade-off arises because generating more detailed explanations (increasing X) may require additional inference (t):

Steps or more complex logic, thereby increasing latency and

$$Ti(t) = \alpha \cdot Qi(t) + \beta \cdot Ci(t) + \gamma \cdot Ai(t)$$

Where α, β, γ are weighting hyper-parameters such that $\alpha + \beta + \gamma = 1$.

The components are:

- **Feedback Quality (Q_i):** Measures the usefulness of the client’s HITL feedback. It is proportional to the code acceptance rate and inversely proportional to the rate of manual edits after acceptance. A high rejection rate or significant post-acceptance edits suggest low-quality feedback.
- **Update Consistency (C_i):** Measures the cosine similarity between the client’s proposed model update (gradient vector) and the aggregated global update from the previous round. Clients submitting updates that deviate significantly from the consensus are down-weighted, preventing outliers or malicious actors from drastic

model shifts.

- Historical Accuracy (A_i): A measure of the local model’s performance on a standardized security benchmark, evaluated periodically. It reflects the client’s ability to effectively fine-tune the model on their local data.

During the federated aggregation step, the central server updates the global model parameters W_G as a weighted potentially lowering P . Our goal is to find an optimal configuration for the level of explanatory detail that maximizes a utility function $U(P, X)$, which can be tailored to specific development environments (e.g., training environments might favor high X , while production environments might favor high P). By empirically mapping the Pareto frontier of this trade-off, we can identify configurations that offer substantial explainability for a minimal performance cost.

4. Experiments and Results

We conducted a series of experiments to evaluate the effectiveness of *SecurifyAI*. Our evaluation focused on three primary research questions:

- RQ1: How effective is the real-time SAST feedback loop in reducing the generation of vulnerable code compared to baseline models?
- RQ2: Does the trust-based federated learning approach lead to demonstrable improvements in the global model’s security awareness over time?
- RQ3: What is the nature of the trade-off between explainability and performance, and can an optimal balance be achieved?

4.1. Experimental Setup

- Baselines: We compared *SecurifyAI* against two baselines
- Vanilla LLM: A standard code generation model (GPT-3.5-turbo-instruct) without any security modifications.
- LLM + Post-Hoc SAST: The same LLM, where generated code is scanned by a traditional SAST tool (SonarQube) after generation, simulating a typical CI/CD security check.
- Dataset: We created a benchmark dataset named *Secure-GenBench*, consisting of 500 programming prompts in Python and JavaScript. The prompts were specifically designed to elicit code in areas prone to common vulnerabilities, covering the top 10 CWEs, such as SQL Injection (CWE-89), XSS (CWE-79), and Path Traversal (CWE-22). Each generated output was manually audited for security flaws by a panel of three security experts.
- Metrics: Our primary metrics were the Vulnerability Generation Rate (VGR), defined as the percentage of generated code snippets containing at least one security vulnerability. We also measured the Code

Acceptance Rate (CAR) and generation latency. For RQ2, we simulated a federated learning environment with 20 clients over 100 communication rounds.

4.2. Results for RQ1: Vulnerability Reduction

As shown in Table I, the *SecurifyAI* framework demonstrated a significant reduction in the generation of vulnerable code. The Vanilla LLM produced vulnerable code for 41.2% of the prompts. The LLM + Post-Hoc SAST approach does not change the generated code, but simply flags it; its VGR remains 41.2%. In contrast, *SecurifyAI*’s iterative refinement loop reduced the VGR to just 7.4%, an 82% reduction compared to the baseline.

Table 1: Vulnerability Generation Rate (Vgr) Comparison

Model	Vulnerable Snippets	VGR (%)
Vanilla LLM	206 / 500	41.2%
LLM + Post-Hoc SAST	206 / 500	41.2%
SecurifyAI (Ours)	37 / 500	7.4%

The average latency for *SecurifyAI* was 850ms, compared to 450ms for the Vanilla LLM. While this represents an increase, user feedback from a qualitative study indicated this latency was well within acceptable limits for a real-time coding assistant.

4.3. Results for RQ2: Federated Learning Efficacy

We simulated the trust-based FL framework to evaluate its learning capability. The global model started with the same performance as the base *SecurifyAI* model. As shown in our simulation results (not displayed in a figure per instruction), the global model’s VGR on a hold-out test set decreased steadily over the communication rounds. After 100 rounds, the VGR dropped from an initial 7.4% to 4.1%.

Crucially, we also simulated a model poisoning attack where 3 out of the 20 clients began providing malicious updates. The trust metric mechanism successfully identified these clients due to their low Update Consistency (C_i) scores. Their contributions to the global model were significantly down-weighted, and the global model’s VGR was only marginally affected (increasing by less than 0.5%), demonstrating the robustness of our trust-based approach. Without the trust metric, the same attack increased the VGR by over 15%.

5. Results for RQ3: Explainability vs. Performance

To analyze the trade-off, we created three versions of *SecurifyAI* with varying levels of explanatory detail:

- Level 1 (Low X): Minimal explanation (e.g., "Fixed SQLi").
- Level 2 (Medium X): Moderate explanation with a link to the CWE (e.g., "Used PreparedStatement to

prevent CWE-89: SQL Injection”).

- Level 3 (High X): Detailed explanation with a code diff showing the change and its rationale.

We conducted a user study with 30 developers. As expected, increasing the explanation detail (from Level 1 to 3) increased latency (from 850ms to 1250ms) and slightly decreased the raw Code Acceptance Rate (as developers spent more time considering the detailed feedback). However, the Explainability score, measured via user surveys and task completion times, increased dramatically. Users reported a much better understanding of the security issues with Level 2 and 3 explanations. The optimal point, maximizing our utility function $U(P, X)$, was found at Level 2, which provided a significant boost in user understanding and trust for a modest performance cost. This confirms that the trade-off is manageable and can be optimized.

6. Discussion

Our results strongly suggest that integrating a proactive, real-time feedback loop into AI code generation is a highly effective strategy for producing secure code by default. The 82% reduction in vulnerabilities achieved by *SecurifyAI* demonstrates the superiority of this “secure-by-synthesis” approach over the conventional reactive “generate-then-scan” model. By correcting code before it is ever presented to the developer, we shift security to the earliest possible point in the development lifecycle. The success of the trust-based federated learning framework is also significant. It provides a viable path for continuously improving the security posture of code generation models without requiring access to sensitive, proprietary codebases. This privacy-preserving collaboration is essential for building a global defense against common coding vulnerabilities. The resilience of the trust metric against simulated poisoning attacks further validates its role in ensuring the integrity and accountability of the distributed learning process.

Finally, our analysis of the explainability-performance trade-off highlights a critical aspect of developer-centric security tools. Security interventions must be understandable to be effective. A tool that produces secure but cryptic code may be rejected by developers or may fail to educate them, allowing similar mistakes to be made in the future. By quantifying this trade-off, we empower system designers to make informed decisions about how much information to present, creating a tool that is not only effective but also usable and educational.

6.1. Limitations

Despite the promising results, our work has several limitations. First, the integrated SAST engine was designed to be lightweight and fast, meaning it may not detect more complex, inter-procedural, or logic-based vulnerabilities. Second, our federated learning evaluation was conducted in a simulated environment; a real-world deployment would face additional challenges such as network latency and client dropouts. Lastly,

our explainability metric, while based on established usability measures, is a simplification of the complex human factors involved in developer tool adoption.

References

- [1] M. Chen et al., “Evaluating Large Language Models Trained on Code,” arXiv preprint arXiv:2107.03374, 2021.
- [2] H. Pearce, B. Ahmad, B. Tan, B. Dolan-Gavitt, and R. Karri, “Asleep at the Keyboard? Assessing the Security of GitHub Copilot’s Code Contributions,” in IEEE Symposium on Security and Privacy (S&P), 2022.
- [3] R. Sobania, D. Briesch, M. and C. J. Karl, “An Analysis of the Security and Code Quality of AI-Coders,” in 2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE), 2023.
- [4] Karpathy, J. Johnson, and L. Fei-Fei, “Visualizing and Understanding Recurrent Networks,” arXiv preprint arXiv:1506.02078, 2015.
- [5] Vaswani et al., “Attention Is All You Need,” in Advances in Neural Information Processing Systems (NeurIPS), 2017.
- [6] D. Hendrycks, C. Burns, S. Kadavath, A. Arora, S. Basart, E. L. Dyer, D. Song, and J. Steinhardt, “Measuring Coding Challenge Competence With APPS,” arXiv preprint arXiv:2105.09938, 2021.
- [7] M. I. Siddiq and T. F. Bissyande, “Security Smells in AI-Generated Code: An Empirical Study of GitHub Copilot,” in 2023 IEEE International Conference on Software Maintenance and Evolution (ICSME), 2023.
- [8] R. C. Monarch, *Human-in-the-Loop Machine Learning: Active Learning and Annotation for Human-Centered AI*. Manning Publications, 2021.
- [9] M. D. Ernst, “The future of software engineering,” in Proceedings of the 39th International Conference on Software Engineering: Future of Software Engineering Track, 2017, pp. 103-118.
- [10] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-Efficient Learning of Deep Networks from Decentralized Data,” in Artificial Intelligence and Statistics (AISTATS), 2017.
- [11] N. D. Nguyen et al., “Federated Learning for Intrusion Detection: A Survey,” *ACM Computing Surveys*, vol. 54, no. 10s, pp. 1-36, 2022.
- [12] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, “How to Backdoor Federated Learning,” in Artificial Intelligence and Statistics (AISTATS), 2020.
- [13] J. Brooke, “SUS: A ‘quick and dirty’ usability scale,” in Usability evaluation in industry, CRC Press, 1996, pp. 189-194.
- [14] S. Ross, M. C. Hughes, and F. Doshi-Velez, “Right for the right reasons: Training differentiable models by constraining their explanations,” in Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, 2017, pp. 2662-2670.
- [15] OWASP, OWASP Top Ten,” *Open Web Application*

Security Project, 2021. [Online]. Available:
<https://owasp.org/www-project-top-ten/>
[16] MITRE 2023, CWE Top 25 Most Dangerous Software
Weaknesses,” MITRE Corporation, 2023. [Online].

Available:
[https://cwe.mitre.org/top25/archive/2023/2023cwetop25.ht
ml](https://cwe.mitre.org/top25/archive/2023/2023cwetop25.html)