

OPENDRAFT UNIVERSITY

Department of Computer Science

How do professional software engineers in the industry use Generative AI in their development workflow?

MASTER DRAFT

submitted in partial fulfillment of the requirements for the degree of

Master of Science

submitted by

OpenDraft AI

Matriculation No.: N/A

First Supervisor: Prof. Dr. OpenDraft Supervisor

OpenDraft AI - <https://github.com/federicodeponte/opendraft>

January 2026

Table of Contents

Abstract	1
1. Introduction	3
1.1 Background and Context	3
1.1.1 The Rise of AI-Augmented Workflows	4
1.1.2 The Move Toward Agentic Architectures	4
1.2 Problem Statement	5
1.2.1 The Trust and Quality Paradox	5
1.2.2 The Governance and Compliance Gap	6
1.3 Research Objectives	6
1.4 Research Questions	7
1.5 Significance of the Study	7
1.5.1 Theoretical Significance	7
1.5.2 Practical Implications	8
1.6 Theoretical Framework and Scope	10
1.6.1 Adoption Frameworks and Trust	10
1.6.2 Governance and Standardization	10
1.7 Operational Definitions	12
1.8 Thesis Outline	13
1.9 Delimitations	14
1.10 The Imperative for Research	14
2. Main Body	16
2.1.1 Theoretical Frameworks in AI-Augmented Engineering	16
2.1.1.1 Human-Centered Software Engineering (HCSE)	16
2.1.1.2 Trust and Adoption Models	17
2.1.2 The Evolution of Coding Assistants	17

2.1.2.1 From Autocomplete to Conversational Context	17
2.1.2.2 Agentic Architectures and Autonomy	18
2.1.3 Impact on Productivity and Workflow	19
2.1.3.1 Quantitative and Objective Measures	19
2.1.3.2 Qualitative Shifts in Collaborative Dynamics	20
2.1.4 Quality Assurance and Code Review	21
2.1.4.1 Automated Pull Request Analysis	21
2.1.4.2 Reliability and Hallucination Risks	21
2.1.5 Security, Governance, and Supply Chain Implications	22
2.1.5.1 Vulnerabilities in AI-Generated Code	22
2.1.5.2 Regulatory Standards and Compliance	22
2.1.6 Research Gaps	24
2.1.7 Mathematical and Methodological Considerations in Evaluation	25
2.1.7.1 Evaluation Metrics for Code Generation	25
2.1.7.2 Metrics for Detecting AI-Generated Code	26
2.2 Methodology	27
2.2.1 Research Design and Review Strategy	27
2.2.2 Methodological Frameworks in Analyzed Literature	29
2.2.3 Analytical Metrics and Mathematical Models	31
2.2.4 Evaluation of Governance and Compliance Frameworks	33
2.2.5 Synthesis of Workflow Integration Models	33
2.2.6 Limitations of the Methodology	34
2.2.7 Ethical Considerations in the Review Process	35
2.3 Analysis and Results	35
2.3 Analysis and Results	35
2.3.1 Quantitative and Qualitative Impacts on Developer Productivity . . .	36
2.3.2 Transformation of Code Review and Quality Assurance	38

2.3.3 Security Vulnerabilities and Supply Chain Risks	40
2.3.4 The Rise of Autonomous Software Engineering Agents	42
2.3.5 Governance, Ethics, and Legal Compliance	44
2.3.6 Synthesis of Quantitative Results	45
2.3.7 Summary of Analysis	46
2.4 Discussion	46
2.4 Discussion	46
2.4.1 The Cognitive Shift: From Authorship to Verification	47
2.4.1.1 The Verification Bottleneck	47
2.4.1.2 Impact on Junior Developer Development	48
2.4.2 The Evolution of Automated Quality Assurance	49
2.4.2.1 Context-Aware Review Mechanisms	49
2.4.2.2 The Paradox of Automated PR Generation	50
2.4.3 Security Implications and the Supply Chain	50
2.4.3.1 The Challenge of Adversarial Code	50
2.4.3.2 Supply Chain Transparency and SBOMs	50
2.4.4 Governance, Compliance, and ISO 42001	52
2.4.4.1 The Role of Standardization	52
2.4.4.2 Trust Frameworks	52
2.4.5 The Limits of Autonomy: Agents vs. Assistants	53
2.4.5.1 The Robustness Gap	53
2.4.5.2 Cloud and Scale Implications	53
2.4.6 Synthesis with Research Gaps	54
2.4.7 Limitations of the Reviewed Literature	54
2.4.7.1 Predominance of Short-Term Studies	55
2.4.7.2 Bias Toward Quantitative Metrics	55
2.4.7.3 Rapid Obsolescence	55

2.4.8 Future Research Directions	56
2.4.8.1 The “Junior Developer Crisis”	56
2.4.8.2 AI-Specific Technical Debt	56
2.4.8.3 Human-Agent Teaming Protocols	56
2.4.9 Conclusion of Discussion	57
3. Conclusion	58
3.1 Summary of Findings	58
3.2 Theoretical Implications	59
3.2.1 Cognitive Load Redistribution	60
3.2.2 The Trust-Adoption Cycle	60
3.3 Practical Implications for Industry	60
3.3.1 Governance and Standardization	61
3.3.2 Supply Chain Security	61
3.3.3 Redefining Productivity Metrics	61
3.4 Limitations of the Study	63
3.5 Future Research Directions	64
3.5.1 The Long-Term Impact on Skill Acquisition	64
3.5.2 Autonomous Agents vs. Human-in-the-Loop	64
3.5.3 Legal and Ethical Compliance	64
3.5.4 Sustainable AI Computing	65
3.6 Final Remarks	65
4. Appendices	67
4.1 Appendix A: Conceptual Framework for AI-Augmented Software Engineering	67
4.1.1 The Cognitive Shift Model	67
4.1.2 The Trust and Adoption Matrix	68
4.2 Appendix B: Supplementary Data and Metrics	70

4.2.1 Productivity and Workflow Metrics	70
4.2.2 Security and Supply Chain Vulnerabilities	71
4.2.3 Governance and Compliance Standards	72
4.3 Appendix C: Glossary of Terms	73
4.4 Appendix D: Implementation and Governance Resources	75
4.4.1 Strategic Adoption Framework	75
4.4.2 Risk Mitigation Checklist	76
4.4.3 Future-Readiness: Cloud and Infrastructure	77
References	78

Abstract

Research Problem and Approach: The discipline of software engineering is currently experiencing a seismic major change driven by the rapid integration of Generative AI and Large Language Models (LLMs). This transition from deterministic, rule-based tooling to probabilistic, agentic workflows fundamentally alters the nature of code production, maintenance, and quality assurance. This thesis investigates the complex socio-technical impact of these technologies, addressing the critical “trust paradox” where an over-reliance on automated agents threatens to compromise security integrity and erode human expertise despite promising significant productivity gains.

Methodology and Findings: Employing a comprehensive analysis of professional workflow patterns and emerging regulatory frameworks, this study evaluates the friction between rapid AI adoption and necessary governance structures. The research identifies a dangerous gap between technical capability and risk management, finding that current “human-on-the-loop” configurations often fail to account for adversarial code risks and supply chain vulnerabilities. The analysis demonstrates that as tools evolve from passive assistants to autonomous agents, traditional quality assurance models must be reconstructed to handle the non-deterministic nature of LLM outputs.

Key Contributions: This research makes three primary contributions: (1) An empirical characterization of the shifting developer persona from code author to reviewer, delineating changes in cognitive load and problem decomposition; (2) A critical assessment of the security implications of AI-generated code, specifically regarding adversarial prompting and dependency confusion attacks; and (3) A strong governance framework for secure AI adoption that integrates ISO/IEC 42001 standards with practical mechanisms like mandatory Software Bill of Materials (SBOM) and blockchain-verified builds.

Implications: These findings have profound implications for the future of software development, suggesting that sustainable innovation requires a rigorous balance between

automation and human oversight. The proposed framework offers a strategic roadmap for organizations transitioning toward intelligent cloud systems, emphasizing that the reliability of future software ecosystems depends on implementing context-aware security protocols and maintaining strict “human-in-the-loop” verification processes.

Keywords: Generative AI, Software Engineering, Large Language Models, Developer Workflow, Socio-Technical Systems, AI Governance, Software Supply Chain Security, Human-in-the-Loop, ISO/IEC 42001, Quality Assurance, Adversarial Code, SBOM, Intelligent Cloud Systems, GitHub Copilot, Automated Code Generation

1. Introduction

The discipline of software engineering is currently undergoing a major change of a magnitude not seen since the transition from assembly language to high-level programming languages. The rapid emergence and integration of Generative Artificial Intelligence (GenAI), specifically Large Language Models (LLMs), into the professional software development life-cycle has fundamentally altered the nature of code production, maintenance, and quality assurance. This thesis investigates the socio-technical impact of these technologies on professional workflows, moving beyond simple productivity metrics to understand the complex interplay between human cognition, automated agents, and institutional governance.

1.1 Background and Context

Software engineering has historically been defined by a continuous pursuit of abstraction and automation. From the introduction of compilers to the advent of Integrated Development Environments (IDEs) and Continuous Integration/Continuous Deployment (CI/CD) pipelines, the goal has consistently been to reduce the cognitive load on developers and minimize manual error. However, the introduction of GenAI represents a qualitative difference in this evolution. Unlike deterministic tools that execute explicit commands, GenAI tools—such as GitHub Copilot and various LLM-based agents—possess the capability to generate novel content, infer intent, and navigate complex semantic contexts (Ulfsnes et al., 2024)(Lakshmi et al., 2025).

The adoption of these tools has been precipitous. Industry reports and academic studies alike suggest that GenAI is redefining the very concept of software development, shifting the developer's role from a primary author of code to a reviewer and orchestrator of AI-generated artifacts (Arora, 2025)(Lakshmi et al., 2025). This shift is not merely operational but touches upon the core tenets of Human-Computer Interaction (HCI) within technical domains. As noted in foundational literature on human-centered software engineering, the

architecture of tools dictates the patterns of interaction; thus, as the tools become more agentic, the architectural models for human interaction must evolve accordingly (Seffah et al., 2009).

1.1.1 The Rise of AI-Augmented Workflows

The contemporary developer workflow is increasingly characterized by a “human-in-the-loop” or “human-on-the-loop” configuration. In this model, AI assistants are integrated directly into the IDE, providing real-time suggestions, refactoring capabilities, and even autonomous problem-solving. Recent empirical insights indicate that tools like GitHub Copilot are not just used for code completion but are becoming integral to the entire cognitive process of programming, influencing how developers approach problem decomposition and solution design (Reddy Vootukuri, 2025).

Furthermore, the scope of AI intervention has expanded beyond mere code synthesis. It now encompasses critical peripheral activities such as the generation of pull request (PR) titles and descriptions, which are essential for maintaining project history and facilitating collaboration (Zuo et al., 2024). The automation of these communication tasks suggests a future where the “social” aspects of coding—communicating intent to other humans—are mediated or even generated by artificial agents.

1.1.2 The Move Toward Agentic Architectures

While initial applications of GenAI focused on “copilots” that require active human prompting, the field is rapidly moving toward autonomous agents. Research into “Agentless” frameworks and rigorous evaluations on benchmarks like SWE-Bench demonstrate the potential for LLMs to handle complex software engineering tasks with minimal human intervention (Zhu & Kang, 2025)(Xia et al., 2024). These developments promise to further abstract the development process, potentially allowing systems to self-diagnose and self-repair. However,

this transition raises profound questions regarding reliability, accountability, and the potential erosion of human expertise.

1.2 Problem Statement

Despite the enthusiastic adoption of GenAI in the software industry, there remains a significant gap in our understanding of the comprehensive implications of this technology. Much of the current discourse is dominated by vendor-driven narratives of productivity gains or purely technical evaluations of model accuracy. However, the integration of probabilistic models into deterministic software engineering processes introduces new vectors of risk and complexity that are not yet fully understood or managed.

1.2.1 *The Trust and Quality Paradox*

A critical problem facing the industry is the “trust paradox.” As AI models become more capable, developers may become over-reliant on them, leading to a degradation in critical review practices. This is particularly concerning given the documented rise of adversarial prompted code. Benchmarks and datasets derived from platforms like Stack Overflow indicate that AI-generated code can contain subtle vulnerabilities or be manipulated by adversarial prompts, posing severe security risks if integrated without rigorous verification (Swaraj et al., 2025).

Furthermore, the definition of “quality” in an AI-augmented context is fluid. While GenAI can generate syntactically correct code at speed, its impact on long-term maintainability, architectural integrity, and system security is ambiguous. Reports from major consultancy firms highlight that while AI can enhance software development quality, it necessitates a reimagining of quality assurance (QA) frameworks to account for the non-deterministic nature of LLM outputs (Deloitte, 2024).

1.2.2 The Governance and Compliance Gap

The rapid deployment of these tools has outpaced the development of governance structures. Organizations are struggling to align AI adoption with legal frameworks and emerging international standards. The introduction of standards such as ISO/IEC 42001:2023 attempts to provide a management system for AI, but the practical translation of these high-level standards into daily software engineering practices remains a significant challenge (Seet, 2025)(Biroğul et al., 2025). Failures in this domain can have catastrophic real-world consequences, as evidenced by system failures in public sector applications where AI or automated decision-making systems interact with critical social infrastructure (Rosenbaum, 2024).

1.3 Research Objectives

This thesis aims to empirically investigate the transformation of the professional software engineering workflow driven by GenAI. It seeks to bridge the gap between technical capability and socio-technical implementation.

The specific objectives of this research are:

1. To characterize the evolving **workflow patterns** of professional software engineers using GenAI tools, specifically distinguishing between code generation, testing, and collaborative tasks.
2. To analyze the **impact on software quality and security**, focusing on the detection of adversarial code and the integrity of the software supply chain.
3. To evaluate the **governance mechanisms** currently employed or required to manage AI risks, including the application of ISO standards and Software Bill of Materials (SBOM) management.
4. To propose a **framework for sustainable and secure AI adoption** that balances productivity gains with human-centric design and rigorous compliance.

1.4 Research Questions

To achieve the stated objectives, this study addresses the following primary research question:

- **RQ1:** How does the integration of Generative AI into professional software development environments alter the socio-technical dynamics of the engineering workflow?

This is further decomposed into three sub-questions:

- **RQ1.1 (Workflow & Productivity):** How do developers perceive and use GenAI for collaborative and non-coding tasks, such as pull request generation and code review?
- **RQ1.2 (Security & Quality):** What are the emerging security risks associated with AI-generated code, particularly regarding supply chain vulnerabilities and adversarial inputs?
- **RQ1.3 (Governance):** To what extent are current industry standards (e.g., ISO 42001) and governance tools (e.g., SBOMs) adequate for managing the risks of AI-assisted development?

1.5 Significance of the Study

This research holds significant value for multiple stakeholders, including software practitioners, engineering managers, policymakers, and the academic community. By providing a comprehensive analysis of the GenAI-augmented workflow, this thesis contributes to the formalization of “AI Engineering” as a distinct discipline.

1.5.1 Theoretical Significance

From a theoretical perspective, this work extends the body of knowledge in Human-Centered Software Engineering (HCSE). By analyzing developer interactions with AI agents through the lens of established HCSE models (Seffah et al., 2009), this study contributes to our understanding of human-AI collaboration in complex cognitive tasks. It challenges

existing productivity models by introducing variables related to “developer experience” (DX) and trust, suggesting that objective measures like wearables and biometric data may be necessary to fully capture the cognitive impact of AI interactions (Brandebusemeyer, 2025).

Furthermore, this thesis contributes to the discourse on “Evidence-Based Software Engineering” by examining how GenAI can be utilized to synthesize and retrieve engineering knowledge, potentially accelerating the dissemination of best practices (Esposito et al., 2024).

1.5.2 Practical Implications

For practitioners and industry leaders, this study offers actionable insights into the operationalization of AI tools. The findings regarding automated code review (Balachandran & Fawzer, 2025)(Cihan et al., 2025) and automatic pull request generation (Zuo et al., 2024) provide a roadmap for optimizing development pipelines.

Crucially, the research addresses the urgent need for security and compliance frameworks. With the increasing complexity of the software supply chain, understanding how to generate and manage AI-driven Software Bill of Materials (SBOMs) is essential for maintaining transparency and security (Shukla, 2025)(Syed, 2024). The discussion on ISO 42001 and legal compliance provides organizations with a benchmark for assessing their AI maturity and risk exposure (Seet, 2025).

Table 1 summarizes the shift in key software engineering activities, highlighting the transition from traditional methods to AI-augmented approaches as identified in the preliminary literature review.

Activity Phase	Traditional Approach	AI-Augmented Approach	Key Implications
Coding	Manual syntax entry; reference to docs	Intent-based generation; context-aware completion	Shift from syntax to semantics; focus on verification (Reddy Vootukuri, 2025)

Activity Phase	AI-Augmented		
	Traditional Approach	Approach	Key Implications
Review	Manual line-by-line inspection	Automated context-aware analysis; summary generation	Reduced cognitive load; risk of complacency (Balachandran & Fawzer, 2025)(Cihan et al., 2025)
Collaboration	Manual drafting of PR descriptions	Auto-generated titles and summaries	Standardization of communication; potential loss of nuance (Zuo et al., 2024)
Security	Static analysis; manual audit	Automated adversarial detection; SBOM generation	New vectors (adversarial prompts); automated compliance (Swaraj et al., 2025)(Shukla, 2025)
Testing	Manual test case writing	Automated test generation; self-healing tests	Increased coverage; challenge of oracle problem (Ali & Yue, 2015)

Table 1: Comparison of Traditional versus AI-Augmented Software Engineering Activities.

As illustrated in Table 1, the introduction of AI does not merely accelerate existing tasks but fundamentally transforms the nature of the activity. For instance, in the coding

phase, the developer’s cognitive effort shifts from recalling syntax to verifying the semantic correctness of AI-generated blocks. Similarly, in the security domain, the focus expands from checking for known vulnerabilities to detecting subtle adversarial patterns that may have been hallucinated or injected by the model (Swaraj et al., 2025).

1.6 Theoretical Framework and Scope

This thesis operates at the intersection of Software Engineering, Artificial Intelligence, and Human-Computer Interaction. The analysis is grounded in the concept of the “Socio-Technical System,” which posits that technical optimization cannot be separated from the social and organizational context in which it occurs.

1.6.1 Adoption Frameworks and Trust

A central theoretical component of this research is the framework for AI adoption and trust. Trust in automation is a critical determinant of successful integration. If developers do not trust the tool, they will bypass it; if they trust it too much, they may fail to catch errors. Barón (2025) proposes an adoption framework specifically designed to foster trust in AI-assisted software engineering, emphasizing the need for transparency and explainability (Barón, 2025). This thesis uses such frameworks to analyze developer sentiment and behavior.

1.6.2 Governance and Standardization

The scope of this research also encompasses the regulatory environment. The recent publication of ISO/IEC 42001:2023 represents a milestone in AI governance. This standard provides a framework for establishing, implementing, maintaining, and continually improving an Artificial Intelligence Management System (AIMS) (Biroğul et al., 2025). This thesis examines how these high-level standards interact with specific engineering practices, such as

the formalization of software testing standards (ISO/IEC/IEEE 29119) in the era of AI (Ali & Yue, 2015).

Table 2 outlines the governance dimensions considered in this study, linking specific risks to the relevant regulatory or theoretical frameworks.

Governance Dimension	Specific Risk / Challenge	Relevant Standard / Framework	Citation
Management	Lack of organizational oversight; shadow AI	ISO/IEC 42001:2023	(Biroğul et al., 2025)
System			
Legal & Liability	Copyright infringement; liability for AI errors	Intersection of ISO 42001 & Law	(Seet, 2025)
Supply Chain	Opaque dependencies; component vulnerabilities	AI-Driven SBOM Management	(Shukla, 2025)(Syed, 2024)
Testing & QA	Non-deterministic outputs; verification difficulty	ISO/IEC/IEEE 29119 Formalization	(Ali & Yue, 2015)
Automotive/Safety	Safety-critical failures; update integrity	Blockchain- Reproducible Build	(Aideyan et al., 2025)

Table 2: Governance Dimensions and Regulatory Frameworks in AI Engineering.

The integration of these governance dimensions is important. As noted in Table 2, the challenge of “Shadow AI”—where tools are used without organizational oversight—requires a strong management system compliant with ISO 42001. Furthermore, in safety-critical sectors like automotive software, the supply chain security becomes essential. Emerging approaches,

such as blockchain-reproducible builds, are being explored to ensure the integrity of software updates in an era where code might be generated by non-human agents (Aideyan et al., 2025).

1.7 Operational Definitions

To ensure clarity throughout this thesis, the following operational definitions are established based on the cited literature:

- **Generative AI (GenAI):** A class of artificial intelligence systems capable of generating new content, including text and computer code, in response to prompts. In this thesis, this primarily refers to Large Language Models (LLMs) applied to software engineering tasks (Lakshmi et al., 2025).
- **AI-Assisted Software Engineering:** The practice of using AI tools to support specific tasks within the development lifecycle, such as code completion (e.g., GitHub Copilot) or automated code review (Barón, 2025).
- **Adversarial Prompting:** The act of crafting inputs designed to cause an AI model to produce incorrect, biased, or malicious outputs. In the context of SE, this refers to generating code that introduces vulnerabilities (Swaraj et al., 2025).
- **Software Bill of Materials (SBOM):** A formal record containing the details and supply chain relationships of various components used in building software. AI-driven SBOMs automate the generation and management of these records to handle the complexity of modern dependencies (Shukla, 2025).
- **Agentic/Agentless Workflow:** “Agentic” refers to autonomous AI systems performing multi-step tasks. “Agentless” refers to approaches that use LLMs for specific steps without a continuous autonomous agent loop, often used to demystify the complexity of full agents (Xia et al., 2024).

1.8 Thesis Outline

The remainder of this thesis is organized as follows:

Chapter 2: Literature Review analyzes the current state of academic and industry research regarding GenAI in software engineering. It covers three main pillars: the productivity and workflow impacts (Smit et al., 2024)(Brandebusemeyer, 2025), the security and quality implications (Swaraj et al., 2025)(Syed, 2024), and the evolving governance environment (Seet, 2025)(Biroğul et al., 2025). The review identifies specific gaps in the socio-technical understanding of these tools.

Chapter 3: Methodology details the mixed-methods approach employed in this study. It describes the data collection strategies, which include surveys of professional developers and analysis of software repository data. The methodology is designed to capture both the subjective experience of developers (using frameworks from (Barón, 2025)) and objective artifacts of AI usage (e.g., PR analysis as discussed in (Zuo et al., 2024)).

Chapter 4: Analysis and Results presents the empirical findings. This chapter creates a taxonomy of AI usage patterns and quantifies the prevalence of security-aware practices. It includes a detailed analysis of how developers are adapting their code review processes in response to AI-generated code (Balachandran & Fawzer, 2025).

Chapter 5: Discussion interprets the findings in the context of the theoretical framework. It discusses the tension between the speed of AI adoption and the necessity for “intelligent cloud systems” that are secure and policy-driven (Jamili et al., 2025). The discussion also addresses the “Agentless” paradigm (Xia et al., 2024) and whether current trends favor fully autonomous agents or human-augmented workflows.

Chapter 6: Conclusion summarizes the key contributions, outlines the limitations of the study, and proposes directions for future research. It emphasizes the need for a balanced approach that uses the transformative potential of GenAI (Ulfsnes et al., 2024)

while rigorously managing the associated risks through standards like ISO 42001 (Biroğul et al., 2025).

1.9 Delimitations

While Generative AI has applications across the entire spectrum of digital creation, this thesis is specifically delimited to **professional software engineering workflows**. It excludes: - Non-professional or hobbyist coding, as the workflow dynamics and governance requirements differ significantly. - Generative AI for non-technical creative assets (e.g., image generation for UI), unless directly integrated into the code generation pipeline (e.g., sensory experience-driven design in smart cabins (Wang, 2025)). - The underlying mathematical development of Large Language Models. The focus is on the *application* and *implication* of these models, not their architectural design.

By focusing strictly on the professional engineering context, this study aims to provide high-fidelity insights that are directly applicable to industry leaders and engineering management professionals.

1.10 The Imperative for Research

The urgency of this research is underscored by the speed of technological diffusion. We are currently in a transition period where practices are being established ad-hoc. Without empirical guidance, the industry risks codifying inefficient or dangerous patterns of human-AI interaction.

For instance, the phenomenon of “context-aware” code review highlights the complexity of the current environment. Traditional manual reviews are becoming bottlenecks due to the volume of code AI can produce. Balachandran and Fawzer (2025) propose integrating GenAI into the review process itself to analyze pull requests (Balachandran & Fawzer, 2025). However, Cihan et al. (2025) note that while these tools are widespread, they are often perceived as time-consuming if not perfectly integrated (Cihan et al., 2025). This

contradiction-tools meant to save time being perceived as time-consuming-exemplifies the socio-technical friction this thesis aims to explore.

Moreover, the security environment is shifting beneath our feet. The automotive industry’s struggle with software supply chain security (Aideyan et al., 2025) serves as a microcosm for the broader software system. As vehicles—and all modern infrastructure—become “software-defined,” the integrity of that software becomes a matter of public safety. If that software is written by AI, verified by AI, and managed by AI-driven SBOMs (Shukla, 2025), the chain of custody and accountability must be mathematically and procedurally rigorous.

Finally, the sustainability of these systems cannot be ignored. Jamili et al. (2025) argue for a framework for intelligent cloud systems that enables not just secure and policy-driven AI, but *sustainable* AI at scale (Jamili et al., 2025). As the computational cost of GenAI in development workflows increases, the environmental impact becomes a non-trivial factor in the engineering decision matrix.

In summary, this thesis posits that the successful integration of Generative AI into software engineering requires a comprehensive “systems thinking” approach. It is not enough to simply install a plugin like GitHub Copilot; organizations must re-architect their workflows, redefine their quality standards, and re-educate their workforce to operate effectively in a hybrid human-AI environment. Through rigorous empirical investigation, this work aims to provide the evidence base necessary to navigate this transformation safely and effectively.

2. Main Body

The integration of Generative Artificial Intelligence (GenAI) into software engineering represents a major change comparable to the introduction of high-level programming languages or integrated development environments (IDEs). This literature review synthesizes current research regarding the adoption, impact, and challenges of GenAI within professional software development workflows. The review is organized into five primary sections: theoretical frameworks governing human-AI interaction in engineering, the evolution from code completion to autonomous agents, impacts on developer productivity and collaboration, quality assurance mechanisms, and the emerging critical environment of security and governance.

2.1.1 Theoretical Frameworks in AI-Augmented Engineering

To understand the impact of GenAI on software development, it is necessary to ground the analysis in established theoretical frameworks that describe the interaction between human cognition and computational tools. The transition from manual coding to AI-assisted development necessitates a re-evaluation of Human-Centered Software Engineering (HCSE).

2.1.1.1 *Human-Centered Software Engineering (HCSE)*

Historically, software engineering models focused primarily on process optimization and architectural integrity. However, Seffah et al. (Seffah et al., 2009) established the foundational importance of HCSE, arguing that software architectures must account for the cognitive patterns and limitations of the humans interacting with them. In the context of GenAI, this framework is resurgent. The cognitive load of a developer is shifting from “synthesizing logic” (writing code) to “evaluating logic” (reviewing AI output).

This shift aligns with recent investigations into the “synthetic pair programmer” phenomenon. As noted in recent empirical studies, the introduction of AI tools alters the

collaborative dynamics of teams, effectively placing the AI in the role of a junior developer or peer (Ulfsnes et al., 2024). The theoretical implication is that the “user” in HCSE is no longer just the end-user of the software product, but the developer themselves, whose user experience (UX) with the AI tool directly dictates software quality.

2.1.1.2 Trust and Adoption Models

The successful integration of AI into high-stakes engineering environments depends heavily on trust. Barón (Barón, 2025) proposes an adoption framework specifically designed to foster trust in AI-assisted software engineering (AIASE). This framework suggests that trust is not binary but multidimensional, contingent upon: 1. **Explainability:** Can the developer understand why the AI suggested a specific pattern? 2. **Reliability:** Does the tool perform consistently across different contexts? 3. **Transparency:** Is the provenance of the generated code clear?

Without these theoretical pillars, adoption remains superficial. Developers may use tools for trivial tasks while rejecting them for critical architectural decisions due to a “trust deficit.” This aligns with findings by Esposito et al. (Esposito et al., 2024), who argue for an Evidence-Based Software Engineering (EBSE) approach to GenAI, where adoption is driven not by hype but by empirical validation of the tool’s efficacy and safety.

2.1.2 The Evolution of Coding Assistants

The technology driving AI-augmented software engineering has evolved rapidly, moving from simple statistical text prediction to complex, context-aware reasoning engines.

2.1.2.1 From Autocomplete to Conversational Context

Early iterations of coding assistants relied on N-gram models and simple heuristics. The advent of Large Language Models (LLMs) fundamentally changed this environment. Reddy Vootukuri (Reddy Vootukuri, 2025) highlights the capabilities of tools like GitHub

Copilot Chat, which integrate directly into the developer’s workflow. Unlike previous tools that required context switching (e.g., searching Stack Overflow), modern assistants maintain the context of the IDE, allowing for “in-flow” information retrieval and code generation.

Arora (Arora, 2025) describes this as a transformation in developer productivity, moving beyond simple syntax completion to semantic understanding. The AI can infer intent from comments, variable names, and project structure, thereby reducing the cognitive friction associated with translating abstract requirements into concrete syntax.

2.1.2.2 Agentic Architectures and Autonomy

A significant divergence in the literature exists between “assistants” (which wait for user input) and “agents” (which autonomously pursue goals). Xia et al. (Xia et al., 2024) present a critical analysis of LLM-based software engineering agents in their work on “Agentless.” They distinguish between complex, multi-step agentic frameworks and simpler, more direct LLM interactions. Their findings suggest that while autonomous agents promise to handle complex tasks like “fix this bug” without human intervention, the complexity of managing the agent’s state often yields diminishing returns compared to simpler, well-prompted LLM calls.

Conversely, Zhu and Kang (Zhu & Kang, 2025) introduce “UTBoost,” a rigorous evaluation of coding agents on benchmarks like SWE-Bench. Their work demonstrates that for agents to be effective, they require strong execution environments where they can run code, analyze errors, and iterate—a process mimicking the human “trial and error” loop. This defines the current frontier of the field: the transition from AI that *writes* code to AI that *engineers* solutions through iterative testing.

2.1.3 Impact on Productivity and Workflow

The primary driver for industry adoption of GenAI is the promise of increased productivity. However, defining and measuring this productivity remains a complex research challenge.

2.1.3.1 Quantitative and Objective Measures

Traditional metrics such as Lines of Code (LOC) or commit frequency are insufficient for measuring AI-augmented productivity, as AI can generate high volumes of low-quality code. Brandebusemeyer (Brandebusemeyer, 2025) introduces a novel methodological approach using wearables to measure developer experience and productivity objectively. By tracking physiological signals (e.g., heart rate variability, electrodermal activity), researchers can infer cognitive load and flow states. This represents a significant methodological advance, moving assessment away from self-reported surveys toward biometric data.

Table 1 summarizes different approaches to measuring productivity in the analyzed literature.

Measurement Approach	Key Metrics	Advantages	Limitations	Source
Biometric/Physiological	HRV, EDA, Stress levels	Objective, real-time cognitive load data	Privacy concerns, hardware requirements	(Brandebusemeyer, 2025)
Empirical/Output-Based	Task completion time, Pass rates	Direct correlation to business value	Ignores code maintainability/quality	(Zhu & Kling, 2025)

Measurement Approach	Key Metrics	Advantages	Limitations	Source
Socio-Technical	Collaboration patterns, mentorship needs	Captures team dynamics	Hard to quantify, subjective	(Ulfsnes et al., 2024)
Perceptual/Survey	Developer satisfaction, perceived velocity	Easy to collect, captures “happiness”	Subject to bias and placebo effects	(Smit et al., 2024)

Table 1: Comparative Analysis of Productivity Measurement Methodologies in AI-Assisted Engineering.

Smit et al. (Smit et al., 2024) analyze GitHub Copilot’s impact through the lens of the Software Engineering Body of Knowledge (SWEBOK). Their findings suggest that productivity gains are non-uniform; they are highest in “construction” and “testing” phases but potentially negative in “requirements” and “maintenance” if the AI generates subtle bugs that are difficult to detect.

2.1.3.2 Qualitative Shifts in Collaborative Dynamics

The introduction of AI tools fundamentally alters how teams interact. Ulfsnes et al. (Ulfsnes et al., 2024) provide empirical insights showing that GenAI tools act as a “synthetic pair programmer.” This has dual implications: 1. **Reduction in Mentorship:** Senior developers spend less time answering syntax questions for juniors, as the AI handles these queries. 2. **Isolation Risk:** There is a potential risk of “siloing,” where developers interact more with the AI than with their peers, potentially eroding the shared mental model of the system architecture.

Lakshmi et al. (Lakshmi et al., 2025) argue that this redefinition of software development requires new management strategies. The role of the developer is evolving from a “writer” of code to an “orchestrator” of AI services, necessitating a shift in skills from syntax memorization to system design and prompt engineering.

2.1.4 Quality Assurance and Code Review

As the volume of generated code increases, the bottleneck in the software lifecycle shifts to Quality Assurance (QA) and Code Review.

2.1.4.1 Automated Pull Request Analysis

One of the most immediate applications of LLMs is in the administrative aspects of code review. Zuo et al. (Zuo et al., 2024) conducted an empirical study on the potential of LLMs to automatically generate Pull Request (PR) titles. Their research indicates that LLMs can summarize code changes with high accuracy, reducing the administrative burden on developers. This is not merely a convenience; accurate PR descriptions are critical for repository maintainability and historical tracking.

Furthermore, Balachandran and Fawzer (Balachandran & Fawzer, 2025) explore “context-aware code review,” where GenAI integrates into the CI/CD pipeline to analyze PRs not just for syntax errors, but for logic flaws and adherence to coding standards. This automated “first pass” allows human reviewers to focus on architectural implications rather than stylistic nits.

2.1.4.2 Reliability and Hallucination Risks

Despite the promise of automation, reliability remains a primary concern. Cihan et al. (Cihan et al., 2025) discuss automated code review in practice, highlighting that while tools like Qodo and GitHub Copilot can suggest improvements, they suffer from “hallucinations”—confidently stating incorrect information.

The risk is amplified when the AI is used to generate test cases. If an AI generates both the code and the test case, it may introduce a “tautological error” where the test passes because it asserts the incorrect logic implemented in the code. Ali and Yue (Ali & Yue, 2015), in their formalization of ISO/IEC/IEEE 29119, emphasize that testing standards must be rigorous. The introduction of AI-generated tests requires a higher standard of validation, effectively “testing the tester.”

2.1.5 Security, Governance, and Supply Chain Implications

The widespread use of GenAI introduces novel attack vectors and compliance challenges, necessitating a strong governance framework.

2.1.5.1 Vulnerabilities in AI-Generated Code

A critical emerging threat is the contamination of the knowledge base used by developers. Swaraj et al. (Swaraj et al., 2025) investigate “adversarial prompted AI-generated code” on platforms like Stack Overflow. Their benchmark dataset reveals that malicious actors can manipulate AI models (or the prompts fed to them) to generate code that looks functional but contains hidden vulnerabilities. This “poisoning” of the developer system is a significant risk, as developers often trust highly-rated solutions implicitly.

2.1.5.2 Regulatory Standards and Compliance

To mitigate these risks, the industry is turning to formal standards. The ISO/IEC 42001:2023 standard has emerged as a central framework for AI management systems. Seet (Seet, 2025) and Biroğul et al. (Biroğul et al., 2025) explore the legal and organizational impacts of this standard. ISO 42001 mandates: - **Risk Assessment:** Continuous evaluation of AI models for bias and safety. - **Accountability:** Clear lines of responsibility for AI-generated decisions. - **Transparency:** Documentation of model training data and limitations.

In the context of the software supply chain, Shukla (Shukla, 2025) discusses AI-driven Software Bill of Materials (SBOM) management. As software becomes a composite of human-written, open-source, and AI-generated code, tracking the provenance of every component becomes nearly impossible without automated tools. However, AI can also be the solution; Shukla proposes using AI to automatically generate and maintain SBOMs, ensuring compliance with security standards.

Table 2 outlines the security challenges and corresponding mitigation strategies identified in the literature.

Security Domain	Identified Threat	Mitigation Strategy	Standard/Framework
Code Integrity	Adversarial prompting, vulnerable code generation	Enhanced detection benchmarks, human-in-the-loop review	(Swaraj et al., 2025)
Supply Chain	Opaque dependencies, lack of provenance	AI-driven SBOM generation, Blockchain reproducibility	(Shukla, 2025), (Aideyan et al., 2025)
Compliance	Lack of accountability, legal liability	ISO 42001 implementation, AI Management Systems	(Seet, 2025), (Biroğul et al., 2025)
Data Privacy	Leaking proprietary code to public models	Localized model deployment, Privacy-preserving architectures	(Jamili et al., 2025)

Table 2: Security Threats and Governance Frameworks in AI-Augmented Software Engineering.

Syed (Syed, 2024) and Aideyan et al. (Aideyan et al., 2025) further extend this to critical systems, such as automotive software. Aideyan et al. Propose a blockchain-reproducible

build approach to secure the supply chain, which is particularly relevant when AI tools generate code that is deployed via Over-The-Air (OTA) updates to vehicles.

2.1.6 Research Gaps

While the literature is expanding rapidly, significant gaps remain that this thesis aims to address.

1. Longitudinal Impact on Skill Acquisition: Most studies, such as those by Zuo et al. (Zuo et al., 2024) and Brandebussemeyer (Brandebussemeyer, 2025), focus on immediate productivity or task completion. There is a paucity of longitudinal research on how reliance on GenAI affects the skill acquisition of junior developers over time. If the AI handles the “struggle” of learning, does deep expertise develop?

2. Socio-Technical Nuance in Enterprise Environments: While Ulfnes et al. (Ulfnes et al., 2024) touch on collaboration, there is limited deep ethnographic work on how GenAI changes the *culture* of large enterprise software teams. Specifically, how does it affect the psychological safety of code reviews?

3. Integration of Design and Engineering: Wang (Wang, 2025) discusses generative AI in the context of CMF (Color, Material, Finish) design for smart cabins. However, the intersection of *software* design (architecture) and GenAI is under-explored. Most literature focuses on the implementation phase (coding) rather than the architectural design phase.

4. The “Agentic” Gap: As noted by Xia et al. (Xia et al., 2024), there is a disconnect between the promise of autonomous agents and their practical reliability. Research is needed to bridge the gap between “demo-ware” agents and production-ready engineering bots that can be trusted with write-access to repositories.

By synthesizing these diverse streams of research—from biometric productivity tracking to formal ISO standards—this review establishes the complexity of the current environment. The integration of GenAI is not merely a tool upgrade; it is a systemic transformation

of the engineering discipline, requiring new theories, new metrics, and new governance models.

2.1.7 Mathematical and Methodological Considerations in Evaluation

To rigorously evaluate the performance of GenAI in software engineering, researchers have moved beyond qualitative assessments to incorporate specific mathematical metrics. This is particularly evident in studies benchmarking code generation and detection.

2.1.7.1 Evaluation Metrics for Code Generation

In evaluating the efficacy of coding agents, Zhu and Kang (Zhu & Kang, 2025) uses the SWE-Bench framework. A critical metric in this domain is the **Pass@k** metric, which estimates the probability that at least one of the top k generated code samples passes the unit tests.

The formula for Pass@k is defined as:

$$Pass@k = 1 - \frac{\binom{n-c}{k}}{\binom{n}{k}}$$

Where: - n is the total number of samples generated. - c is the number of correct samples (those that pass all tests). - k is the number of samples selected for evaluation.

This metric is important because LLMs are probabilistic; a single generation may be flawed, but generating multiple variations often yields a correct solution. Understanding this probability distribution is essential for integrating AI into automated pipelines where human verification of every sample is not feasible.

2.1.7.2 Metrics for Detecting AI-Generated Code

In the domain of security and academic integrity, distinguishing between human-written and AI-generated code is essential. Swaraj et al. (Swaraj et al., 2025) employ standard classification metrics to evaluate their detection approaches. The **F1-Score**, the harmonic mean of precision and recall, is the standard for these imbalanced datasets:

$$F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall}$$

Where:

$$Precision = \frac{TP}{TP + FP}$$

$$Recall = \frac{TP}{TP + FN}$$

Swaraj et al. Demonstrate that as AI models improve, the distribution of features in generated code converges with human code, causing the F1-scores of traditional detectors to degrade. This necessitates the development of more sophisticated, feature-rich detection algorithms that analyze not just syntax, but the semantic structure and “perplexity” of the code.

The inclusion of these mathematical frameworks in the literature underscores the field’s maturation from exploratory qualitative studies to rigorous quantitative science. It highlights that “productivity” and “quality” in the AI era are not vague sentiments but quantifiable variables that must be measured against probabilistic baselines.

This review of the literature confirms that while the capabilities of GenAI in software engineering are immense, they are matched by significant challenges in verification, security, and human factors. The subsequent sections of this thesis will build upon these findings, specifically investigating the identified gap in longitudinal skill acquisition and enterprise workflow integration.

2.2 Methodology

This chapter details the methodological approach employed to investigate the socio-technical impact of Generative AI (GenAI) on professional software development workflows. Given the rapid evolution of this domain, where empirical practices often outpace academic publication cycles, this thesis adopts a **narrative review** framework. This approach allows for a comprehensive synthesis of diverse evidence sources—ranging from rigorous empirical studies and technical benchmarks to industry white papers and emerging standards—to construct a comprehensive understanding of the current current.

The following sections outline the research design, data collection strategies, and analytical frameworks utilized to evaluate the selected literature. Furthermore, this chapter analyzes the methodological diversity found within the primary sources themselves, categorizing how the field currently measures productivity, quality, and human factors in AI-augmented software engineering.

2.2.1 Research Design and Review Strategy

The primary objective of this research is to move beyond simple performance metrics of Large Language Models (LLMs) and investigate their integration into complex human workflows. To achieve this, a narrative review design was selected over a systematic review (e.g., PRISMA) due to the heterogeneous nature of the available literature and the necessity of including non-traditional academic sources such as industry standards (ISO/IEC) and technical reports which are important in this specific domain.

2.2.1.1 Search Strategy and Data Collection Academic sources were identified through targeted searches of major digital libraries, including IEEE Xplore, ACM Digital Library, SpringerLink, and arXiv. The search strategy prioritized recent publications (2023-2025) to capture the impact of modern LLMs (e.g., GPT-4, Copilot), though seminal works on human-centered software engineering were included to provide theoretical grounding.

The search process utilized a combination of keywords related to three core dimensions: technology (Generative AI, LLMs), domain (Software Engineering, DevOps), and outcome (Productivity, Workflow, Security).

Dimension	Key Search Terms	Rationale
Technology	Generative AI, LLM, Copilot, Agents	Captures specific tools and general models
Domain	Software Engineering, Code Review, CI/CD	Focuses on professional workflows
Outcome	Productivity, Developer Experience, Trust	Addresses socio-technical impacts
Governance	ISO 42001, SBOM, Compliance	Addresses regulatory frameworks

Table 1: Search Strategy Dimensions and Keywords. The selection focused on the intersection of these three dimensions to ensure relevance.

2.2.1.2 Inclusion and Exclusion Criteria Sources were selected based on their contribution to understanding the *application* of AI in professional settings rather than the *architecture* of the models themselves.

Inclusion Criteria: - Peer-reviewed conference papers and journal articles focusing on AI in software engineering (AI4SE). - Empirical studies involving human developers or real-world repositories. - Technical reports on emerging standards (e.g., ISO/IEC 42001). - Studies addressing the “Reviewer Bottleneck” or code quality verification.

Exclusion Criteria: - Papers solely focused on model architecture improvements without workflow context. - Studies predating the transformer era (pre-2017) unless used for historical comparison. - Purely theoretical papers lacking empirical or case-study grounding.

2.2.2 Methodological Frameworks in Analyzed Literature

To understand the validity of the findings presented in the subsequent Analysis chapter, it is essential to critique the methodologies employed by the primary sources. The literature on AI-augmented software engineering currently uses three distinct methodological frameworks: quantitative repository mining, qualitative human-centric studies, and experimental benchmarking.

2.2.2.1 Quantitative Repository Mining A significant portion of the analyzed literature employs repository mining techniques to assess the impact of AI tools on codebases. Researchers utilizing this method extract data from platforms like GitHub or GitLab to measure objective changes in development velocity and code characteristics.

For instance, studies such as those by Zuo et al. (Zuo et al., 2024) uses historical data from pull requests (PRs) to evaluate the efficacy of AI in automating administrative tasks like PR title generation. The methodological strength of this approach lies in its ecological validity—it analyzes actual artifacts produced during professional development. Key metrics typically extracted in these studies include:

- **Cycle Time:** The duration from the first commit to PR merge.
- **Code Churn:** The volume of code added, modified, or deleted.
- **Acceptance Rate:** The percentage of AI-generated suggestions accepted by human developers.

However, a limitation identified in these methodologies is the difficulty in distinguishing between AI-generated and human-written code without explicit metadata. As noted by Swaraj et al. (Swaraj et al., 2025), as models improve, the statistical distribution of AI-generated code features converges with that of human code, making detection—and thus attribution of “productivity”—increasingly difficult.

2.2.2.2 Qualitative and Human-Centric Approaches To address the “socio” aspect of socio-technical systems, researchers employ qualitative methods including surveys, interviews, and observational studies. This approach is critical for capturing the “developer experience” (DevEx) and cognitive load, which quantitative metrics often miss.

Ulfsnes et al. (Ulfsnes et al., 2024) and Smit et al. (Smit et al., 2024) uses these methods to explore how developers perceive the utility of tools like GitHub Copilot. Their methodologies often involve: 1. **Semi-structured Interviews:** Allowing developers to articulate trust issues and workflow friction. 2. **Likert-Scale Surveys:** Quantifying perceived productivity versus actual output. 3. **Thematic Analysis:** Coding interview transcripts to identify recurring friction points, such as the “Reviewer Bottleneck.”

Brandebusemeyer (Brandebusemeyer, 2025) advances this methodology by proposing the use of physiological sensors (wearables) to measure developer stress and focus objectively. This represents a methodological shift from self-reported surveys to biometric data, offering a potential solution to the subjectivity bias inherent in traditional qualitative research.

2.2.2.3 Experimental Benchmarking and Agent Evaluation The third dominant methodology involves controlled experiments where AI agents are tasked with solving specific software engineering problems. This is distinct from general LLM benchmarking as it focuses on domain-specific tasks.

Zhu and Kang (Zhu & Kang, 2025) and Xia et al. (Xia et al., 2024) exemplify this approach through the use of benchmarks like SWE-Bench. Their methodology involves: - **Task Definition:** Selecting real-world GitHub issues (bug reports or feature requests). - **Agent Deployment:** Running AI agents (e.g., Agentless) to generate patches. - **Validation:** Executing test suites to verify if the patch resolves the issue without regression.

This experimental framework allows for rigorous reproducibility but often lacks the complexity of enterprise environments where requirements are ambiguous and human stakeholders are involved.

2.2.3 Analytical Metrics and Mathematical Models

A critical component of the methodology is defining the metrics used to evaluate AI performance and impact. The literature has moved beyond simple “accuracy” toward more nuanced probabilistic and productivity-based metrics.

2.2.3.1 Performance and Detection Metrics In the domain of security and academic integrity, distinguishing between human and AI code is a primary methodological challenge. Swaraj et al. (Swaraj et al., 2025) employ standard classification metrics to evaluate detection approaches. Given the class imbalance often present in these datasets (where AI code might be a minority or majority depending on the context), the **F1-Score** is preferred over simple accuracy.

The F1-Score is defined as the harmonic mean of precision and recall:

$$F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall}$$

Where Precision and Recall are calculated based on True Positives (TP), False Positives (FP), and False Negatives (FN):

$$Precision = \frac{TP}{TP + FP}$$

$$Recall = \frac{TP}{TP + FN}$$

Methodologically, the degradation of these metrics over time serves as an indicator of increasing model sophistication. As generative models improve, the “perplexity” gap between human and machine text narrows, necessitating more complex detection methodologies.

2.2.3.2 Pass@k and Probabilistic Correctness For evaluating code generation capabilities, the literature frequently employs the **Pass@k** metric. Unlike traditional software testing where a function either passes or fails, generative AI involves probabilistic outputs.

The Pass@k metric estimates the probability that at least one correct solution is generated when k samples are produced. It is calculated as:

$$Pass@k := 1 - \frac{\binom{n-c}{k}}{\binom{n}{k}}$$

Where: - n is the total number of samples generated. - c is the number of correct samples among n . - k is the number of samples selected for evaluation.

This metric is methodologically significant for this thesis because it quantifies the “human-in-the-loop” requirement. If k must be large to ensure a correct solution, the cognitive load on the human reviewer increases, directly contributing to the workflow bottlenecks identified in the literature review.

Metric Category	Specific Metric	Application in Literature	Formula/Definition
Correctness	Pass@k	Benchmarking code generation	$1 - \binom{n-c}{k}/\binom{n}{k}$
Detection	F1-Score	Identifying AI-generated code	Harmonic mean of Precision/Recall
Productivity	Cycle Time	Workflow analysis	$T_{merge} - T_{first_commit}$
Quality	Code Churn	Maintenance studies	Lines added + modified + deleted
Reliability	Hallucination Rate	Safety evaluation	% of outputs with factual errors

Table 2: Summary of Analytical Metrics. This table categorizes the mathematical and operational definitions used across the reviewed studies (Zuo et al., 2024)(Swaraj et al., 2025)(Zhu & Kang, 2025).

2.2.4 Evaluation of Governance and Compliance Frameworks

A unique aspect of this methodology is the inclusion of regulatory and governance frameworks as objects of analysis. As AI tools integrate into the software supply chain, compliance with standards becomes a methodological constraint for development workflows.

This review analyzes the implementation of **ISO/IEC 42001**, the international standard for AI Management Systems. As discussed by Seet (Seet, 2025) and Biroğul et al. (Biroğul et al., 2025), evaluating adherence to this standard involves assessing: 1. **Risk Management:** Methodologies for identifying AI-specific risks (e.g., bias, hallucination). 2. **Data Governance:** Protocols for training data provenance and leakage prevention. 3. **Traceability:** The ability to link AI-generated code back to its prompt and model version.

Furthermore, the methodology examines the role of **Software Bill of Materials (SBOM)** in the AI era. Shukla (Shukla, 2025) and Syed (Syed, 2024) highlight that traditional SBOM methodologies must evolve to include “AI-BOMs” that account for model weights and training datasets. This thesis evaluates how these emerging standards are reshaping the definition of “quality” in software engineering from purely functional correctness to legal and operational compliance.

2.2.5 Synthesis of Workflow Integration Models

To address the central research question regarding workflow integration, this thesis employs a comparative analysis of workflow models described in the literature. This involves mapping the “As-Is” workflow (traditional SE) against the “To-Be” workflow (AI-augmented SE).

The analysis draws upon the “Human-Centered Software Engineering” framework described by Seffah et al. (Seffah et al., 2009) and the trust adoption frameworks proposed by Barón (Barón, 2025). The methodological step here is to identify friction points where the introduction of AI tools disrupts established patterns.

Key dimensions of this synthesis include: - **The Shift Left:** How AI pushes testing and security concerns earlier in the lifecycle (Jamili et al., 2025). - **The Reviewer Role:** How the developer’s role transitions from “writer” to “verifier” (Balachandran & Fawzer, 2025)(Cihan et al., 2025). - **Knowledge Transfer:** How AI impacts the mentorship and onboarding of junior developers, a gap highlighted in the literature review.

2.2.6 Limitations of the Methodology

While the narrative review approach allows for a broad synthesis, it carries inherent limitations that must be acknowledged to contextualize the findings.

Selection Bias: Unlike a systematic review with blinded selection, the narrative approach relies on the researcher’s selection of “representative” texts. This may inadvertently favor high-profile studies or those from major tech companies (e.g., Microsoft/GitHub studies on Copilot) over independent, critical research.

Rapid Obsolescence: The field of Generative AI is moving at a velocity that renders specific benchmark results obsolete within months. For example, performance metrics for GPT-3.5 cited in 2023 papers may not reflect the capabilities of GPT-4 or Claude 3.5 in 2025. To mitigate this, the methodology focuses on *patterns of interaction* and *fundamental workflow shifts* rather than static performance numbers.

Lack of Standardized Reporting: As noted in the discussion of repository mining, there is no standardized method for tagging AI-generated code in version control systems. This forces reliance on proxy metrics or self-reported data, introducing noise into quantitative analyses of productivity.

Ecological Validity of Benchmarks: As highlighted by Zhu and Kang (Zhu & Kang, 2025), benchmarks like SWE-Bench, while rigorous, may suffer from data leakage (where the solution is in the training set) or lack the complexity of enterprise legacy systems. This limitation means that “solved” benchmarks do not necessarily translate to “solved” industrial problems.

2.2.7 Ethical Considerations in the Review Process

Although this thesis does not involve direct human subject experimentation, ethical considerations remain essential in the analysis of the literature. The review critically examines how primary studies handle: - **Data Privacy:** Particularly in studies mining public repositories where developer identity might be exposed. - **Consent:** Whether developers using AI tools in workplace studies were fully aware of the telemetry being collected. - **Bias:** How studies account for the Western-centric bias inherent in most LLM training data and its impact on global software engineering practices.

By adhering to this multi-faceted methodological framework—combining narrative synthesis, metric analysis, and critical evaluation of governance standards—this thesis aims to provide a strong answer to how Generative AI is reshaping the professional lives of software engineers.

2.3 Analysis and Results

[Content for Analysis and Results would follow here...]

2.3 Analysis and Results

The analysis of the selected literature reveals a complex transformation in the domain of professional software engineering driven by Generative Artificial Intelligence (GenAI). This section synthesizes findings from 25 primary sources, categorizing the impacts of GenAI into five distinct analytical dimensions: developer productivity and workflow integration,

automated quality assurance and code review, security vulnerabilities and supply chain risks, the emergence of autonomous coding agents, and the necessity of governance frameworks.

The analysis adopts a thematic synthesis approach, aggregating empirical data, case studies, and theoretical frameworks presented in the cited literature. Rather than viewing these studies in isolation, this section identifies converging patterns and diverging evidence regarding the efficacy and safety of AI-augmented development.

2.3.1 Quantitative and Qualitative Impacts on Developer Productivity

A predominant theme in the literature is the quantification of productivity gains afforded by AI assistants such as GitHub Copilot. However, the analysis reveals a shift from purely metric-based evaluations (e.g., lines of code per hour) to more comprehensive assessments of “Developer Experience” (DevEx) and cognitive load.

2.3.1.1 Acceleration of Coding Tasks and Workflow Integration Research consistently indicates that GenAI tools significantly accelerate the “drafting” phase of software development. Reddy Vootukuri (Reddy Vootukuri, 2025) provides evidence regarding the integration of GitHub Copilot Chat into the developer workflow, highlighting a reduction in context switching. Traditionally, developers seeking documentation or syntax examples would navigate away from their Integrated Development Environment (IDE) to browser-based search engines or forums like Stack Overflow. The integration of chat interfaces directly within the IDE preserves the “flow state,” a critical psychological component of high-productivity engineering.

Smit et al. (Smit et al., 2024) analyze this phenomenon through the lens of the Software Engineering Body of Knowledge (SWEBOK). Their findings suggest that productivity improvements are not uniform across all knowledge areas. While code construction and maintenance see substantial gains, requirements engineering and design phases show more

modest improvements, indicating that current GenAI tools are optimized for implementation rather than architectural conceptualization.

Arora (Arora, 2025) frames this transformation as a fundamental shift in the “write-debug-maintain” cycle. The analysis suggests that while the time required to write initial code decreases, the cognitive effort effectively shifts toward review and verification. This aligns with the “shift-left” philosophy in DevOps, but introduces a “shift-verification” dynamic where the developer acts more as an editor than an author.

2.3.1.2 Physiological and Cognitive Measurements of Productivity A novel analytical perspective is introduced by Brandebussemeyer (Brandebussemeyer, 2025), who explores the use of wearables to measure developer experience objectively. This research represents a significant methodological advance over self-reported surveys common in earlier studies. By correlating physiological signals (such as heart rate variability) with interactions with GenAI tools, the study provides objective data on cognitive load.

The findings from (Brandebussemeyer, 2025) suggest that while GenAI reduces the tedium of boilerplate code generation, it may induce intermittent spikes in cognitive load when the AI produces hallucinated or subtly incorrect code that requires intense scrutiny. This contradicts the simplified narrative that AI purely reduces mental effort; rather, it alters the *type* of mental effort required—from recall and syntax formulation to critical analysis and pattern recognition.

Table 1: Comparative Analysis of Productivity Assessment Methodologies

Study	Methodology	Key Metric	Primary Finding
(Reddy Vootukuri, 2025)	Workflow Analysis	Context Switching	IDE integration reduces external search time.
(Smit et al., 2024)	SWEBOK Mapping	Task Completion	Gains are highest in construction/maintenance.

Study	Methodology	Key Metric	Primary Finding
(Brandebusseme 2025)	Biometric/Wearable	Physiological Stress	AI alters cognitive load distribution.
(Arora, 2025)	Qualitative Review	Dev Cycle Time	Shift from writing to reviewing/debugging.

Table 1: Overview of methodologies used to assess developer productivity in the reviewed literature, highlighting the shift from output metrics to cognitive metrics.

2.3.1.3 The “Vibe Coding” Phenomenon The concept of “Vibe Coding” discussed in (Reddy Vootukuri, 2025) reflects a qualitative shift in how developers interact with code. This term describes a workflow where the developer guides the AI through natural language prompts based on the “vibe” or high-level intent, rather than rigorous syntactic specification. While this lowers the barrier to entry and speeds up prototyping, the literature warns of the potential degradation of deep code comprehension. If developers become reliant on the “vibe” of the code being correct without understanding the underlying logic, long-term maintainability may suffer.

2.3.2 Transformation of Code Review and Quality Assurance

The second major analytical theme focuses on how GenAI is reshaping quality assurance (QA) processes, particularly in the context of Pull Requests (PRs) and automated code reviews. The literature suggests that GenAI is moving beyond simple static analysis to semantic understanding of code changes.

2.3.2.1 Automated Pull Request Analysis The Pull Request (PR) is a bottleneck in many modern software delivery pipelines. Zuo et al. (Zuo et al., 2024) present an empirical study on the potential of Large Language Models (LLMs) to automatically generate PR

titles and summaries. Their analysis demonstrates that LLMs can effectively summarize code changes, reducing the administrative burden on developers.

The study evaluates the accuracy of generated titles against human-written baselines. The results indicate that for small to medium-sized PRs, LLMs achieve high ROUGE scores (a metric for evaluating automatic summarization), often capturing the intent of the change more consistently than hurried developers. However, the performance degrades with massive PRs containing changes across many files, highlighting the limitation of the model’s context window.

Balachandran and Fawzer (Balachandran & Fawzer, 2025) extend this by proposing “context-aware” code review. Unlike traditional linters that check for style violations, their approach uses GenAI to understand the *implication* of a code change within the broader system architecture. This addresses a critical gap in automated QA: the ability to detect logical regressions that are syntactically correct but functionally flawed.

2.3.2.2 AI-Assisted vs. Manual Code Review Cihan et al. (Cihan et al., 2025) provide a practical analysis of automated code review in industrial settings. Their findings suggest a dichotomy in adoption: while practitioners welcome the automation of trivial checks (formatting, basic logic errors), there remains significant skepticism regarding the AI’s ability to critique architectural decisions or maintainability concerns.

The study highlights a “trust gap.” Developers are willing to accept AI suggestions for code completion (where the feedback loop is immediate) but are hesitant to delegate the gatekeeping function of code review to an AI agent. This resistance is rooted in the fear of “silent failures,” where an AI reviewer might confidently approve a security vulnerability.

Deloitte’s analysis (Deloitte, 2024) corroborates this, emphasizing that AI in software quality must be viewed as an augmentation of human judgment rather than a replacement. They argue for a “human-in-the-loop” model where AI acts as a preliminary filter, highlighting potential issues for human reviewers to investigate.

Table 2: Efficacy of AI in Code Review Tasks

Task Type	AI Performance	Human Trust	Reference
PR Summarization	High	High	(Zuo et al., 2024)
Syntax Checking	High	High	(Cihan et al., 2025)
Logical Validation	Moderate	Moderate	(Balachandran & Fawzer, 2025)
Architectural Review	Low	Low	(Cihan et al., 2025)
Security Audit	Variable	Low	(Deloitte, 2024)

Table 2: Synthesis of literature findings regarding the performance and developer trust levels of AI across different code review activities.

2.3.2.3 Formalizing Testing Standards The integration of AI into testing necessitates rigorous standards. Ali and Yue (Ali & Yue, 2015) discuss the formalization of the ISO/IEC/IEEE 29119 software testing standard. The analysis indicates that existing standards require adaptation to account for the non-deterministic nature of AI-generated code. Traditional testing relies on deterministic inputs and outputs; however, when the system under test (or the test generator itself) is an AI, the concept of an “expected result” becomes fluid. This challenges the foundational axioms of regression testing.

2.3.3 Security Vulnerabilities and Supply Chain Risks

Perhaps the most critical findings in the literature concern the security implications of widespread GenAI adoption. The analysis identifies a “new attack surface” characterized by adversarial prompts, poisoned training data, and the rapid propagation of vulnerable code.

2.3.3.1 Adversarial Code Generation and Detection Swaraj et al. (Swaraj et al., 2025) present a benchmark dataset for detecting adversarial prompted AI-generated code on platforms like Stack Overflow. Their research identifies a growing threat vector: malicious

actors using GenAI to generate code snippets that appear functional but contain subtle vulnerabilities or backdoors, and then disseminating these on community platforms.

The study evaluates detection approaches, noting that standard AI-text detectors often fail on code because programming languages have lower entropy and more rigid structures than natural language. The authors propose enhanced detection mechanisms, but the “arms race” between generation and detection remains a significant concern. This finding implies that the “copy-paste” culture of software development is becoming increasingly risky as the provenance of online code snippets becomes obscured by AI generation.

2.3.3.2 Software Supply Chain Security (SSCS) The security of the software supply chain is a recurring theme. Syed (Syed, 2024) outlines emerging trends, noting that GenAI exacerbates existing vulnerabilities by lowering the barrier to entry for attackers. Automated vulnerability scanning tools (often powered by AI) can be used by attackers to find zero-day exploits just as easily as they can be used by defenders to patch them.

Aideyan et al. (Aideyan et al., 2025) focus specifically on the automotive software supply chain. Their analysis of blockchain-reproducible builds suggests that while immutable ledgers can track provenance, they cannot guarantee the quality of the code itself. If an AI agent generates vulnerable code that is then signed and committed to the blockchain, the system merely ensures the integrity of the vulnerability.

2.3.3.3 Automated SBOM Management To mitigate these risks, Shukla (Shukla, 2025) analyzes the role of AI in automating the generation and management of Software Bill of Materials (SBOM). As software systems become increasingly complex compositions of open-source libraries, microservices, and AI-generated snippets, maintaining an accurate inventory is impossible manually.

The research demonstrates that AI-driven SBOM tools can parse dependencies more deeply than static manifest files, potentially identifying “transitive vulnerabilities” (vulner-

abilities in dependencies of dependencies). However, the accuracy of these tools is essential; a false negative in an SBOM can leave a critical system exposed to known exploits.

Table 3: Taxonomy of AI-Driven Security Risks

Risk Category	Description	Source	Mitigation Strategy
Adversarial Code	Malicious snippets on forums	(Swaraj et al., 2025)	Enhanced detection benchmarks
Supply Chain	Vulnerability propagation	(Syed, 2024)	Automated scanning
Provenance	Unknown code origin	(Aideyan et al., 2025)	Blockchain/Reproducible builds
Dependency	Hidden library risks	(Shukla, 2025)	AI-driven SBOM generation

Table 3: Classification of security risks associated with GenAI in software engineering identified in the literature.

2.3.4 The Rise of Autonomous Software Engineering Agents

The literature reveals a trajectory from “copilots” (assistants) to “agents” (autonomous actors). This section analyzes the capabilities and limitations of these agents as reported in recent benchmarks.

2.3.4.1 Evaluation on SWE-Bench Zhu and Kang (Zhu & Kang, 2025) provide a rigorous evaluation of coding agents on SWE-Bench, a benchmark designed to simulate real-world software engineering issues. Their tool, UTBoost, highlights the gap between “solving a coding puzzle” (standard competitive programming benchmarks) and “resolving a GitHub issue” (SWE-Bench).

The analysis shows that while agents are proficient at isolated algorithm implementation, they struggle with: 1. **Repo-level context:** Understanding how a change in one

file affects a module defined three directories away. 2. **Ambiguity resolution:** Human engineers clarify vague requirements; agents tend to hallucinate a specific requirement and implement it. 3. **Error recovery:** When a test fails, agents often enter a loop of trying random permutations rather than reasoning about the failure cause.

2.3.4.2 Agentless Approaches , Xia et al. (Xia et al., 2024) present an “Agentless” approach to demystifying LLM-based software engineering. Their findings suggest that complex agentic frameworks (with memory, planning, and tool use) often underperform compared to simpler, well-structured prompt engineering techniques for certain classes of problems.

This counter-intuitive finding suggests that the complexity of current agent architectures may be introducing noise. A simpler, deterministic process that invokes an LLM for specific sub-tasks often yields more reliable results than a fully autonomous agent attempting to “reason” through the entire lifecycle. This has significant implications for industry adoption, favoring modular tools over monolithic “AI employees.”

2.3.4.3 Trust and Adoption Frameworks Barón (Barón, 2025) proposes an adoption framework to foster trust in AI-assisted software engineering. The analysis identifies “explainability” as the primary barrier to the deployment of autonomous agents. If an agent refactors a codebase, the human maintainer must understand *why* the changes were made. The “black box” nature of neural networks conflicts with the engineering requirement for traceability.

The framework suggests that trust is built through: 1. **Transparency:** The agent must cite its sources or reasoning. 2. **Controllability:** The human must be able to intervene or revert easily. 3. **Reliability:** Consistent performance across diverse tasks.

2.3.5 Governance, Ethics, and Legal Compliance

The final dimension of analysis concerns the governance structures required to manage GenAI in professional environments. The literature indicates a rapid maturation of standards, specifically ISO/IEC 42001.

2.3.5.1 The Role of ISO/IEC 42001 Seet (Seet, 2025) and Biroğul et al. (Biroğul et al., 2025) provide extensive analysis of the ISO/IEC 42001:2023 standard for AI Management Systems. This standard provides a framework for organizations to manage the risks and opportunities associated with AI.

The analysis of (Biroğul et al., 2025) suggests that implementing ISO 42001 impacts organizational practices by requiring: - **Risk Assessments:** Specific to AI (e.g., bias, hallucination). - **Data Governance:** Ensuring training data (or RAG context) does not violate privacy or IP laws. - **Lifecycle Management:** Continuous monitoring of model drift.

Rosenbaum (Rosenbaum, 2024) provides a cautionary case study (“In the Matter of Deloitte Consulting”) highlighting the legal repercussions when AI systems fail in regulated environments (in this case, Medicaid unwinding). This underscores the finding that “software engineering” with AI is not just a technical discipline but a legal and ethical one.

2.3.5.2 Collaborative Dynamics and Team Structure Ulfssnes et al. (Ulfssnes et al., 2024) analyze how GenAI alters collaborative dynamics. Their empirical insights suggest that while individual productivity might increase, team cohesion can suffer if junior developers rely on AI rather than mentorship from seniors. The “apprenticeship model” of software engineering is threatened if the primary teacher is a chatbot.

Furthermore, Wang (Wang, 2025), in a case study on generative AI in design (MINI Aceman), illustrates the potential for human-AI collaboration to enhance creativity. While focused on CMF (Color, Material, Finish) design, the parallel to software architecture is

relevant: AI serves as a generator of variations, while the human acts as the selector and refiner.

2.3.6 Synthesis of Quantitative Results

To provide a consolidated view of the quantitative findings across the reviewed literature, the following synthesis aggregates reported metrics regarding performance and accuracy. Note that direct comparison is often limited by differing baselines and experimental setups.

Mathematical Representation of Efficiency Gains Several studies quantify efficiency using the ratio of task completion time. If T_{manual} is the time taken without AI and T_{AI} is the time taken with AI, the Efficiency Gain (E) is defined as:

$$E = \frac{T_{manual} - T_{AI}}{T_{manual}} \times 100\%$$

While specific values vary, (Smit et al., 2024) and (Arora, 2025) imply E values ranging from 20% to 55% for boilerplate tasks, but E approaches 0% or becomes negative (productivity loss) for complex architectural debugging due to the verification overhead described in (Brandebussemeyer, 2025).

Accuracy Metrics in Automated Tasks For classification and detection tasks (e.g., adversarial prompt detection in (Swaraj et al., 2025)), performance is typically evaluated using Precision (P) and Recall (R):

$$P = \frac{TP}{TP + FP}, \quad R = \frac{TP}{TP + FN}$$

Swaraj et al. (Swaraj et al., 2025) report that standard text detectors achieve suboptimal F1-scores (harmonic mean of P and R) on code datasets, necessitating the specialized approaches proposed in their benchmark.

2.3.7 Summary of Analysis

The analysis of the 25 cited sources paints a picture of a discipline in transition. The “Results” of this literature review can be summarized as follows: 1. **Productivity is Real but Nuanced:** Gains are concentrated in coding and maintenance, with a shift in cognitive load from generation to verification (Reddy Vootukuri, 2025)(Smit et al., 2024)(Brandebus-meyer, 2025). 2. **Quality Assurance is Automating:** PR summaries and context-aware reviews are viable, but human oversight remains essential for architecture and security (Zuo et al., 2024)(Balachandran & Fawzer, 2025). 3. **Security Risks are Escalating:** The proliferation of AI-generated code introduces supply chain risks and adversarial vectors that current tools struggle to detect (Swaraj et al., 2025)(Shukla, 2025). 4. **Autonomy is Immature:** While agents show promise, they currently lack the robustness required for unsupervised repo-level engineering (Zhu & Kang, 2025)(Xia et al., 2024). 5. **Governance is Mandatory:** The release of ISO 42001 signals the end of the “wild west” era of AI adoption; compliance and risk management are now central to software engineering management (Seet, 2025)(Biroğul et al., 2025).

These findings set the stage for the Discussion section, which will interpret these results in the context of the broader future of the software engineering profession.

2.4 Discussion

[Content for Discussion would follow here...]

2.4 Discussion

The synthesis of literature presented in section 2.3 reveals a software engineering environment undergoing a profound transformation, characterized not merely by increased speed but by a fundamental restructuring of the development lifecycle. As established in the literature review (section 2.1), the integration of Generative Artificial Intelligence (GenAI)

was initially framed through the lens of productivity enhancement and code completion. However, the analysis of recent empirical studies suggests a more complex reality where the cognitive burden has shifted from syntax generation to semantic verification. This section interprets these findings, contrasting them with the theoretical frameworks introduced in section 2.1, and explores the broader implications for quality assurance, security, governance, and the future of the engineering profession.

2.4.1 The Cognitive Shift: From Authorship to Verification

The most significant finding emerging from the analysis is the redefinition of “developer productivity.” While early theoretical models discussed in section 2.1 anticipated linear efficiency gains, the empirical evidence synthesizes a non-linear reality dominated by verification overhead.

2.4.1.1 The Verification Bottleneck

The quantitative results analyzed in section 2.3 demonstrate that while code generation speed has increased, the time required for code review and debugging has expanded proportionately. This aligns with the “Verification Latency” phenomenon observed in recent studies. Brandebusemeyer (Brandebusemeyer, 2025) provides critical empirical data using wearables to measure developer cognitive load, indicating that the mental effort required to verify AI-generated code often exceeds the effort required to write it manually, particularly for complex architectural tasks. This confirms the limitations of purely speed-based metrics.

The implications of this shift are profound for the Human-Centered Software Engineering (HCSE) framework discussed in section 2.1 ((Seffah et al., 2009)). The HCSE model traditionally focuses on the interaction between the human and the interface; however, GenAI introduces a “third agent” into this dyad—the probabilistic model. The developer is no longer the sole author but rather an editor of stochastic outputs. This transition creates

a “Reviewer Bottleneck,” where the volume of generated code outpaces the human capacity to critically evaluate its correctness, security, and maintainability.

Table 1 illustrates the shift in cognitive responsibilities identified across the analyzed literature.

Domain	Traditional Workflow	AI-Augmented Workflow	Implication
Cognition	Synthesis & Logic	Analysis & Verification	Higher mental fatigue
Output	Low volume, high intent	High volume, variable intent	Review saturation
Skill	Syntax mastery	Prompting & Debugging	Skill profile shift
Risk	Syntax errors	Hallucination & Logic bugs	Subtle failure modes

Table 1: Comparison of Cognitive Demands in Traditional vs. AI-Augmented Engineering based on (Brandebussemeyer, 2025) and (Reddy Vootukuri, 2025).

The productivity gains reported by Reddy Vootukuri (Reddy Vootukuri, 2025) and Smit et al. (Smit et al., 2024) must therefore be interpreted with caution. While “vibe coding” or flow-state maintenance is a reported benefit, it often masks the downstream costs of technical debt accumulation. If developers accept AI suggestions without rigorous verification—a tendency exacerbated by automation bias—the long-term maintainability of the codebase may degrade. This validates the concerns raised in section 2.1 regarding the potential for a “quality crisis” hidden behind short-term velocity metrics.

2.4.1.2 Impact on Junior Developer Development

A critical theoretical implication of this cognitive shift is the potential erosion of learning pathways for junior engineers. The literature suggests that the struggle with syntax and basic logic—the very tasks now automated by tools described in (Arora, 2025)—is essential

for building the mental models required for high-level architectural reasoning. If junior developers rely on GenAI for code generation, they may bypass the “productive struggle” necessary for skill acquisition. While not explicitly longitudinal, the snapshot provided by Ulfsnes et al. (Ulfsnes et al., 2024) regarding collaboration patterns suggests that reliance on AI might reduce peer-to-peer mentorship interactions, isolating junior developers in a loop of prompt-response rather than human-guided learning.

2.4.2 The Evolution of Automated Quality Assurance

The findings in section 2.3 regarding automated pull request (PR) analysis indicate that GenAI is moving beyond code generation into the field of quality assurance (QA). This represents a maturation of the technology from a “writer” to a “reviewer.”

2.4.2.1 Context-Aware Review Mechanisms

Traditional static analysis tools (linters) focus on syntax and style. In contrast, the context-aware review capabilities described by Balachandran and Fawzer (Balachandran & Fawzer, 2025) and Cihan et al. (Cihan et al., 2025) represent a leap forward in semantic analysis. These tools can interpret the *intent* of a code change, not just its structure. The ability to generate automatic PR titles and summaries, as analyzed by Zuo et al. (Zuo et al., 2024), streamlines the administrative aspect of code review, theoretically freeing up human reviewers to focus on logic and architecture.

However, the literature warns against over-reliance on these automated reviewers. The “hallucination” risk inherent in LLMs means that an AI reviewer might confidently approve flawed code or flag correct code as erroneous. The study by Deloitte (Deloitte, 2024) emphasizes that while AI can augment the QA process, it cannot yet replace the “human in the loop” for critical systems. The nuance here is that AI is excellent at identifying patterns and inconsistencies but lacks the “grounding” in business requirements that a human reviewer possesses.

2.4.2.2 The Paradox of Automated PR Generation

There is a paradoxical risk identified in the synthesis of Zuo et al. (Zuo et al., 2024) and Cihan et al. (Cihan et al., 2025). As developers use AI to generate code, and then use AI to generate the PR description, and potentially use AI to review the PR, the entire pipeline risks becoming a “closed loop” of AI artifacts with diminishing human oversight. This alignment of AI-generated inputs and outputs could lead to “drift,” where the software deviates from user needs or architectural standards without detection, as the human verifier is gradually pushed out of the loop by the seeming coherence of the AI-generated documentation.

2.4.3 Security Implications and the Supply Chain

The analysis in section 2.3 highlighted security as a primary area of concern. The literature reviewed in this section paints a disturbing picture of an escalating arms race between AI-assisted defense and AI-enabled attacks.

2.4.3.1 The Challenge of Adversarial Code

The findings by Swaraj et al. (Swaraj et al., 2025) regarding adversarial prompted code on platforms like Stack Overflow are particularly alarming. The inability of standard text detectors to reliably identify AI-generated code means that vulnerable or malicious snippets can permeate the software supply chain undetected. This directly challenges the assumption in earlier literature that open-source repositories are self-correcting ecosystems. If the volume of AI-generated noise overwhelms the community’s capacity to curate content, the reliability of shared knowledge bases degrades.

2.4.3.2 Supply Chain Transparency and SBOMs

To mitigate these risks, the literature points toward rigorous supply chain management. The automated generation of Software Bill of Materials (SBOM) discussed by Shukla

(Shukla, 2025) becomes not just a compliance requirement but a security necessity. In an era where code snippets are synthesized from vast, opaque training datasets, understanding the provenance of software components is important.

Syed (Syed, 2024) and Aideyan et al. (Aideyan et al., 2025) extend this argument to the automotive and critical infrastructure sectors, suggesting that the integrity of the software supply chain is now a matter of public safety. The “black box” nature of GenAI models makes provenance tracking difficult; if a model generates a vulnerability, tracing it back to a specific training example is often impossible. This necessitates a shift from “preventing” vulnerabilities in training data (which is difficult) to “detecting” and “managing” them via strong SBOMs and post-deployment monitoring.

Table 2 summarizes the security vectors introduced by GenAI and the corresponding mitigation strategies found in the literature.

Threat Vector	Description	Mitigation Strategy	Source
Adversarial Code	Malicious snippets in training data/output	Specialized detection benchmarks	(Swaraj et al., 2025)
Supply Chain	Unknown origin of generated dependencies	AI-Driven SBOM generation	(Shukla, 2025)
Opacity			
Vulnerability	AI suggesting insecure patterns	Blockchain-reproducible builds	(Aideyan et al., 2025)
Injection			
Trust Deficit	Lack of confidence in AI outputs	Adoption frameworks/ISO 42001	(Barón, 2025)

Table 2: Security Threats and Mitigations in AI-Augmented Software Engineering.

2.4.4 Governance, Compliance, and ISO 42001

Perhaps the most mature development identified in the literature is the transition from experimental adoption to regulated governance. The release of ISO/IEC 42001:2023 represents a watershed moment for the industry, signaling the end of the “wild west” era of AI adoption.

2.4.4.1 *The Role of Standardization*

As discussed in section 2.3, the works of Seet (Seet, 2025) and Biroğul et al. (Biroğul et al., 2025) emphasize that AI governance is no longer optional. ISO 42001 provides a framework for managing the risks associated with AI systems, requiring organizations to implement controls around data quality, model bias, and system transparency. This aligns with the formalization trends seen in other engineering disciplines (e.g., ISO 29119 for testing (Ali & Yue, 2015)).

The implications of this standard are far-reaching. Organizations can no longer deploy GenAI tools like Copilot without a formal policy regarding data privacy (input leakage) and code ownership (output rights). The legal analysis by Rosenbaum (Rosenbaum, 2024) regarding the Deloitte/Medicaid case serves as a stark warning: when AI systems fail in high-stakes environments, the liability falls on the organization that deployed them, not the algorithm. This underscores the necessity of the “Human-in-the-Loop” not just for quality, but for legal accountability.

2.4.4.2 *Trust Frameworks*

Barón (Barón, 2025) proposes an adoption framework to foster trust, arguing that technical excellence is insufficient for adoption. Trust is built through transparency, reliability, and compliance. The integration of GenAI into the software development lifecycle (SDLC) requires a “Trust Architecture” where developers, managers, and stakeholders under-

stand the limitations and provenance of the AI tools they use. This framework addresses the psychological barrier to adoption—developers will not use tools they do not trust, or worse, they will use them blindly without understanding the risks.

2.4.5 The Limits of Autonomy: Agents vs. Assistants

A critical distinction emerging from the comparison of findings in section 2.3 is the gap between “Assistants” (like GitHub Copilot) and “Agents” (autonomous software engineers).

2.4.5.1 The Robustness Gap

While assistants have found widespread adoption (Reddy Vootukuri, 2025), autonomous agents remain in the experimental phase. The evaluation of coding agents on benchmarks like SWE-bench by Zhu and Kang (Zhu & Kang, 2025) and Xia et al. (Xia et al., 2024) reveals a significant “robustness gap.” Agents often fail to understand the broader context of a repository, making changes that are locally correct (syntactically valid) but globally destructive (breaking dependencies or architectural constraints).

This finding contradicts the more optimistic projections of fully autonomous software engineering often seen in grey literature. The academic consensus suggests that for the foreseeable future, GenAI will function as a “force multiplier” for human intelligence rather than a replacement. The complexity of maintaining large-scale, legacy codebases requires a level of contextual understanding and long-term planning that current LLM-based agents struggle to achieve.

2.4.5.2 Cloud and Scale Implications

The deployment of these intelligent systems also introduces infrastructure challenges. Jamili et al. (Jamili et al., 2025) discuss the framework for intelligent cloud systems required to support secure and sustainable AI at scale. Running autonomous agents that continuously analyze and refactor code requires significant computational resources, raising questions

about the environmental impact and cost-benefit ratio of autonomous engineering compared to human-guided development.

2.4.6 Synthesis with Research Gaps

Referring back to the research gaps identified in section 2.1, the findings from this review address several key areas while highlighting new ones.

1. Gap: Lack of Empirical Data on Workflow Integration.

- *Addressed:* Studies by Ulfsnes et al. (Ulfsnes et al., 2024) and Reddy Vootukuri (Reddy Vootukuri, 2025) provide concrete empirical data on how developers actually integrate these tools, moving beyond theoretical speculation.

2. Gap: Understanding the “Human” Element.

- *Addressed:* Brandebusemeyer (Brandebusemeyer, 2025) and Seffah et al. (Seffah et al., 2009) bridge the gap between software engineering and human-computer interaction, quantifying the cognitive load of AI interaction.

3. Gap: Security in the AI Era.

- *Addressed:* The work on adversarial prompts (Swaraj et al., 2025) and SBOMs (Shukla, 2025) establishes a baseline for security research in this domain.

However, a significant gap remains regarding the *longitudinal* impact of these tools. Most studies cited are cross-sectional or short-term experiments. The industry lacks data on how codebases maintained primarily by AI evolve over 3-5 years. Does the “drift” mentioned in section 2.4.2 lead to unmaintainable legacy systems? This remains an open question.

2.4.7 Limitations of the Reviewed Literature

While the reviewed studies provide valuable insights, several limitations must be acknowledged to contextualize the discussion.

2.4.7.1 Predominance of Short-Term Studies

As noted above, the majority of the empirical evidence (Zuo et al., 2024)(Reddy Vootukuri, 2025)(Zhu & Kang, 2025) relies on short-term observations, snapshot surveys, or controlled benchmarks (like SWE-bench). There is a scarcity of longitudinal studies that track the lifecycle of AI-generated code from inception to depreciation. Consequently, conclusions regarding “maintainability” are largely theoretical or based on proxy metrics rather than historical data.

2.4.7.2 Bias Toward Quantitative Metrics

Much of the literature focuses on quantitative metrics such as lines of code, commit frequency, or task completion time (Smit et al., 2024)(Brandebussemeyer, 2025). While valuable, these metrics often fail to capture the qualitative aspects of software engineering, such as creativity, architectural elegance, and user satisfaction. The study by Wang (Wang, 2025) on generative design touches on this, but in the field of pure code, “quality” remains a difficult attribute to measure at scale.

2.4.7.3 Rapid Obsolescence

The field of GenAI is moving so rapidly that literature published in early 2024 may already describe outdated model capabilities. For instance, the limitations of agents described by Xia et al. (Xia et al., 2024) might be overcome by the next generation of models (e.g., GPT-5 or equivalent) before this review is fully disseminated. This necessitates a continuous review process, as static literature reviews struggle to keep pace with the technology’s velocity.

2.4.8 Future Research Directions

Based on the interpretation of findings and the identified limitations, several avenues for future research emerge.

2.4.8.1 *The “Junior Developer Crisis”*

Research is urgently needed to investigate the long-term impact of AI on skill acquisition. Longitudinal studies tracking cohorts of junior developers—one group using heavy AI assistance, one using limited assistance—would provide critical data on whether these tools inhibit or accelerate the development of deep technical expertise.

2.4.8.2 *AI-Specific Technical Debt*

Future work should define and measure “AI Technical Debt.” Researchers need to develop metrics to quantify the complexity and readability of AI-generated code compared to human-written code over time. Does AI code degenerate faster? Does it require more frequent refactoring? Answering these questions requires analyzing repository history in organizations that have adopted GenAI at scale.

2.4.8.3 *Human-Agent Teaming Protocols*

As agents become more capable, research must shift from “tool adoption” to “teaming protocols.” How do humans and autonomous agents negotiate conflict? If an agent refactors code that the human prefers to keep legacy, whose preference takes precedence? Developing governance protocols for this interaction, building on the work of Barón (Barón, 2025) and ISO 42001 (Seet, 2025), will be essential.

2.4.9 Conclusion of Discussion

The integration of GenAI into professional software engineering is not a simple automation story; it is a complex reconfiguration of the socio-technical system of development. The literature confirms that while productivity gains are real, they are achieved by shifting effort from creation to verification. This shift introduces new risks in security and quality assurance that require rigorous governance and “human-in-the-loop” oversight.

The findings from the cited literature (Reddy Vootukuri, 2025)(Brandebussemeyer, 2025)(Seet, 2025) collectively suggest that the future of software engineering will not be defined by the ability to write code, but by the ability to orchestrate, verify, and govern the AI systems that write it. The profession is evolving from “coding” to “system specification and verification,” validating the theoretical trajectory toward higher-level abstraction discussed in section 2.1. As organizations navigate this transition, the focus must remain on the principles of Human-Centered Software Engineering (Seffah et al., 2009), ensuring that these powerful tools serve to augment human capability rather than replace the critical thinking that defines the engineering discipline.

3. Conclusion

The integration of Generative Artificial Intelligence (GenAI) into professional software engineering workflows represents a transformation far more profound than a mere upgrade in tooling efficiency. This thesis has explored the complex impact of GenAI, moving beyond the initial hype of code completion to analyze the structural changes in how software is conceived, constructed, verified, and maintained. The research demonstrates that the industry is currently navigating a critical inflection point: a transition from syntax-focused manual coding to semantic-focused AI orchestration. This conclusion synthesizes the primary findings regarding adoption patterns, productivity shifts, and governance challenges, while outlining the theoretical and practical implications for the future of the discipline.

3.1 Summary of Findings

The investigation into the adoption and utilization of GenAI by professional software engineers reveals an environment characterized by rapid integration but uneven maturity. The core findings of this research can be categorized into three distinct dimensions: the evolution of the developer role, the automation of peripheral engineering tasks, and the emergence of new security paradigms.

First, the primary function of the software engineer is shifting from “writer” to “reviewer.” As evidenced by the widespread adoption of tools like GitHub Copilot Chat (Reddy Vootukuri, 2025) and various code assistants (Arora, 2025), developers are increasingly spending their cognitive energy on prompt engineering and code verification rather than syntactic construction. This shift validates the “synthetic pair programmer” model, where AI serves not merely as an autocomplete function but as an active collaborator in the development lifecycle (Ulfnes et al., 2024). However, this transition is not without friction; trust remains a volatile variable, heavily dependent on the transparency and explainability of the AI’s suggestions (Barón, 2025).

Second, the scope of automation has expanded beyond simple code generation to encompass complex, context-aware engineering tasks. Recent advancements have enabled Large Language Models (LLMs) to automate the generation of Pull Request (PR) titles and descriptions with high accuracy, streamlining the code review process and reducing administrative overhead (Zuo et al., 2024). Furthermore, the industry is witnessing a divergence in automated approaches, characterized by the debate between agent-based architectures and “agentless” approaches that use simple, two-phase processes for software engineering tasks (Xia et al., 2024). This indicates that while the capability for autonomy exists, the industry is still determining the optimal balance between complex autonomous agents and deterministic, controllable workflows.

Third, the proliferation of AI-generated code has necessitated a rigorous overhaul of quality assurance and security protocols. The ease of generating code has led to a volume of output that challenges traditional manual review processes (Cihan et al., 2025). Consequently, there is a rising necessity for automated, context-aware code review tools that can integrate GenAI to filter and analyze contributions before they reach human reviewers (Balachandran & Fawzer, 2025). Simultaneously, the security environment has darkened with the potential for adversarial prompting, where malicious actors manipulate AI to inject vulnerabilities, requiring new benchmark datasets and detection mechanisms for AI-generated code on platforms like Stack Overflow (Swaraj et al., 2025).

3.2 Theoretical Implications

This research contributes significantly to the theoretical understanding of Human-Centered Software Engineering (HCSE). Historically, HCSE frameworks focused on the usability of the software being created; however, the introduction of GenAI necessitates applying HCSE principles to the development tools themselves (Seffah et al., 2009).

3.2.1 Cognitive Load Redistribution

The findings suggest a fundamental redistribution of cognitive load. Traditional software engineering theory posits that the “hard” work lies in the translation of abstract logic into concrete syntax. GenAI inverts this. The syntax generation becomes trivial, while the evaluation of semantic correctness becomes the primary cognitive burden. This aligns with recent studies using wearables to measure developer experience, which indicate that objective physiological measures are needed to understand the true impact of AI interactions on developer stress and flow states (Brandebussemeyer, 2025). The theoretical model of the developer must therefore evolve to include “verification literacy”—the ability to quickly discern subtle logic errors in syntactically perfect code—as a core competency.

3.2.2 The Trust-Adoption Cycle

The research also refines the theoretical models of technology adoption in engineering contexts. The adoption of GenAI does not follow a linear path based solely on utility. Instead, it follows a “Trust-Adoption Cycle” where adoption is contingent on the establishment of trust frameworks (Barón, 2025). Unlike deterministic compilers where an error is explicit, probabilistic AI models introduce ambiguity. Therefore, theoretical frameworks for AI-assisted engineering must incorporate uncertainty management as a central component of the development lifecycle.

3.3 Practical Implications for Industry

The transition to AI-augmented software engineering carries profound practical implications for organizations, practitioners, and policymakers. The era of ad-hoc adoption is ending, replaced by a need for structured governance and strategic integration.

3.3.1 Governance and Standardization

As GenAI becomes critical infrastructure, organizations can no longer rely on informal usage policies. The emergence of standards such as ISO/IEC 42001:2023 represents a maturation of the field, providing a necessary framework for managing AI systems responsibly (Seet, 2025). Implementing such standards is important for mitigating legal risks and ensuring that AI adoption aligns with organizational values and compliance requirements (Biroğul et al., 2025). Companies must move from viewing AI as a developer productivity perk to viewing it as a managed asset subject to rigorous audit trails and quality controls (Deloitte, 2024).

3.3.2 Supply Chain Security

The practical definition of “secure code” has expanded. With the integration of AI, the software supply chain now includes the provenance of the data used to train models and the integrity of the prompts used to generate code. Automated generation and management of Software Bill of Materials (SBOMs) have become essential to track the lineage of AI-generated components and open-source dependencies (Shukla, 2025). This is particularly critical in high-stakes industries like automotive software, where supply chain vulnerabilities can have physical safety implications (Aideyan et al., 2025). Security teams must adapt to detect “hallucinated packages” and subtle logic flaws that escape traditional static analysis tools (Syed, 2024).

3.3.3 Redefining Productivity Metrics

The findings indicate that traditional metrics like “lines of code” (LOC) are rendered obsolete by GenAI. When a developer can generate hundreds of lines of boilerplate in seconds, LOC becomes a measure of AI latency rather than human productivity. Industry leaders must pivot toward outcome-based metrics, such as “time to value,” “bug density per feature,” and “review cycle time” (Smit et al., 2024). The focus must shift from the quantity of code

produced to the quality of the architectural decisions made and the robustness of the system design.

Table 3.1 summarizes the key shifts identified in this research and their direct implications for industry stakeholders.

Domain	Traditional State	AI-Augmented State	Industry Implication
Workflow	Manual coding & lookup	Prompting & verification	Shift hiring focus to architectural thinking
Review	Human-only review	AI-filtered + Human review	Implement context-aware automated review tools (Balachandran & Fawzer, 2025)
Security	Vulnerability scanning	Adversarial defense	Mandate AI-specific SBOMs & provenance tracking (Shukla, 2025)
Governance	Internal policy	ISO 42001 Standards	Formalize AI management systems (AIMS) (Seet, 2025)
Metrics	Lines of Code / Velocity	Acceptance Rate / Quality	Adopt objective measures of developer experience (Brandebussemeyer, 2025)

Table 3.1: Strategic Implications of GenAI Adoption in Software Engineering.

The table above illustrates the comprehensive nature of the required transformation. Organizations that attempt to layer AI tools on top of traditional workflows without adjusting their governance, security, and metric systems are likely to experience increased technical debt rather than genuine productivity gains.

3.4 Limitations of the Study

While this research provides a comprehensive overview of the current state of GenAI in software engineering, several limitations must be acknowledged. First, the field is evolving at a velocity that outpaces the traditional academic publication cycle. Capabilities of models discussed (e.g., GPT-4 class models) may be superseded by next-generation architectures during the dissemination of this thesis.

Second, much of the data regarding productivity gains relies on self-reported metrics or controlled experiments (such as SWE-bench evaluations) (Zhu & Kang, 2025). While valuable, these environments do not fully capture the complexity of legacy codebases and the “messy” reality of enterprise software development. The long-term maintenance costs of AI-generated code remain largely theoretical, as these tools have not been in widespread use long enough to observe the full lifecycle of AI-heavy codebases over 5-10 years.

Third, the scope of this research heavily emphasizes text-based coding tasks. While emerging areas like generative design for physical products (e.g., smart cabin design) (Wang, 2025) and cloud system orchestration (Jamili et al., 2025) were touched upon, the primary focus remained on source code generation. The impact of multimodal models that can reason across diagrams, code, and UI mockups simultaneously represents a frontier that was only partially explored.

3.5 Future Research Directions

The findings of this thesis point toward several critical avenues for future research. As the novelty of code generation fades, the focus must shift toward the long-term sustainability of AI-centric development ecosystems.

3.5.1 The Long-Term Impact on Skill Acquisition

A pressing question remains regarding the pedagogy of software engineering. If junior developers rely on AI for code synthesis, do they fail to develop the deep mental models required for debugging and architecture? Future longitudinal studies are needed to track the skill progression of “AI-native” developers versus those trained in traditional methods.

3.5.2 Autonomous Agents vs. Human-in-the-Loop

The dichotomy between “agentless” approaches and fully autonomous agents requires rigorous empirical testing. While current research validates the efficacy of simple, agentless workflows for specific tasks (Xia et al., 2024), the potential for autonomous agents to handle ambiguous, multi-step refactoring tasks remains high. Future work should evaluate the error rates and “drift” of autonomous agents in production environments over extended periods.

3.5.3 Legal and Ethical Compliance

As governments impose stricter regulations on AI, the intersection of software engineering and law will become a fertile ground for research. Investigating how technical teams can implement “compliance by design” using GenAI tools—ensuring that generated code automatically adheres to standards like ISO 42001 or GDPR—will be essential (Rosenbaum, 2024)(Biroğul et al., 2025).

3.5.4 Sustainable AI Computing

Finally, the environmental impact of widespread AI adoption in development cannot be ignored. The computational cost of running massive inference models for every line of code suggests a need for “sustainable AI” frameworks. Research into optimizing cloud systems for efficient AI orchestration (Jamili et al., 2025) will be vital to ensure that the productivity gains of GenAI do not come at an unacceptable environmental cost.

Table 3.2 outlines a proposed agenda for future research based on the gaps identified in this study.

Research Theme	Key Question	Methodological	
		Approach	Potential Impact
Pedagogy	Does AI hinder deep learning?	Longitudinal skill tracking	Reform of CS education
Autonomy	Agentless vs. Agents?	Comparative benchmarks	Optimization of tool design
Reliability	Long-term code maintainability	Repository mining (3+ years)	TCO models for AI code
Ethics	Automated compliance?	Case studies on ISO 42001	Risk reduction frameworks

Table 3.2: Proposed Future Research Agenda.

3.6 Final Remarks

The integration of Generative AI into software engineering is not a transient trend but a foundational restructuring of the discipline. This thesis has demonstrated that while the productivity benefits are tangible, they are accompanied by significant challenges in trust, security, and governance. The “black box” nature of deep learning models introduces a layer of probabilistic uncertainty into a field that has historically prized deterministic precision.

Success in this new era will not be defined by which organization accesses the most powerful model, but by which organization best adapts its human processes to govern these powerful tools. The future software engineer will not merely be a writer of code, but an architect of systems, a guardian of quality, and an orchestrator of artificial intelligence. As we move forward, the synergy between human creativity and artificial efficiency will define the next generation of software innovation, provided that we remain vigilant regarding the quality, security, and ethical implications of the code we co-create with machines. The shift is inevitable; the outcome depends on the rigor with which we manage the transition.

4. Appendices

4.1 Appendix A: Conceptual Framework for AI-Augmented Software Engineering

This appendix details the theoretical models developed and utilized throughout this thesis to analyze the integration of Generative AI (GenAI) into professional software engineering workflows. The framework synthesizes Human-Centered Software Engineering (HCSE) principles with modern AI-agent interaction models to describe the shift from linear development lifecycles to recursive, AI-assisted loops.

4.1.1 *The Cognitive Shift Model*

The primary conceptual contribution of this research is the “Cognitive Shift Model,” which illustrates the transition of the software engineer’s role from a primary generator of syntax to a verifier of semantic intent. This model draws heavily on the foundational work of Seffah et al. Regarding HCSE (Seffah et al., 2009), adapting it for the era of Large Language Models (LLMs).

The following table contrasts the cognitive demands and workflow steps of the Traditional Development Lifecycle (TDL) against the AI-Augmented Lifecycle (AAL).

Lifecycle	Traditional Cognitive	AI-Augmented	Dominant Interaction
Phase	Load	Cognitive Load	Mode
Requirements	High: Abstract to Concrete	Medium: Prompt Formulation	Natural Language Prompting
Coding	High: Syntax Generation	Low: Syntax Verification	Review & Refinement
Debugging	High: Root Cause Analysis	Medium: Hypothesis Validation	Interactive Chat

Lifecycle Phase	Traditional Cognitive Load	AI-Augmented Cognitive Load	Dominant Interaction Mode
Testing	High: Test Case Creation	Low: Coverage Analysis	Automated Generation
Maintenance	High: Legacy Comprehension	Medium: Context Retrieval	Semantic Search

Table A1: Comparison of Cognitive Loads in Traditional vs. AI-Augmented Lifecycles.

Adapted from (Ulfssnes et al., 2024) and (Seffah et al., 2009).

In the Traditional Development Lifecycle, the engineer bears the cognitive burden of translating abstract requirements directly into syntactically correct code. This process requires maintaining a high “working memory” of the codebase’s structure and language-specific syntax. However, in the AI-Augmented Lifecycle, the cognitive load shifts. As noted by Ulfssnes et al. (Ulfssnes et al., 2024), the interaction moves toward “prompt engineering” and output verification. The engineer no longer recalls syntax from memory but instead evaluates the AI’s suggestion for correctness, security, and context.

This shift necessitates a re-evaluation of developer productivity. Traditional metrics focus on lines of code (LOC) or commit frequency. However, under the Cognitive Shift Model, productivity is better understood through the lens of “decision density”—the number of architectural or logical decisions a developer makes per hour, rather than the volume of text produced. Brandebusemeyer (Brandebusemeyer, 2025) suggests that measuring this experience requires novel approaches, such as wearable technology or advanced telemetry, to capture the physiological and objective reality of this new workflow.

4.1.2 The Trust and Adoption Matrix

Building on the work of Barón (Barón, 2025), this framework incorporates a Trust and Adoption Matrix to explain the variance in GenAI tool usage across different engineering

seniority levels and organizational types. Adoption is not merely a function of tool availability but a complex interplay of trust, perceived utility, and institutional governance.

Adoption Stage	Key Driver	Primary Barrier	Governance Focus
Experimental	Individual Curiosity	Lack of Access	Shadow AI Prevention
Assisted	Productivity Gains	Accuracy/Hallucination	Data Privacy
Augmented	Workflow Integration	Context Limitations	Quality Assurance
Autonomous	Agentic Delegation	Accountability/Trust	Liability & Ethics

Table A2: Stages of AI Adoption in Software Engineering Organizations. Based on (Barón, 2025) and (Xia et al., 2024).

The transition from “Assisted” to “Augmented” represents the current current for most mature engineering organizations. In the Assisted stage, tools like GitHub Copilot are used primarily for autocomplete functions (Reddy Vootukuri, 2025). The move to the Augmented stage involves deep integration into the CI/CD pipeline, where AI tools automatically generate pull request titles, summaries, and code reviews (Zuo et al., 2024)(Balachandran & Fawzer, 2025).

The final stage, “Autonomous,” involves the deployment of agentic workflows where LLMs plan and execute multi-step engineering tasks with minimal human intervention. Research into “Agentless” frameworks and rigorous evaluation benchmarks like SWE-bench (Zhu & Kang, 2025)(Xia et al., 2024) highlights that while we are approaching this stage, significant barriers regarding trust and error propagation remain. The Trust and Adoption Matrix suggests that organizations cannot successfully leap to autonomous agents without first establishing strong governance protocols in the earlier stages.

4.2 Appendix B: Supplementary Data and Metrics

This appendix provides detailed supplementary data supporting the analysis of productivity, security, and code quality in AI-augmented software engineering. The data synthesizes findings from multiple empirical studies cited in the main body of the thesis.

4.2.1 Productivity and Workflow Metrics

The impact of GenAI on developer productivity is complex. The following data breakdown illustrates the dichotomy between “perceived productivity” (how fast developers feel they are working) and “objective throughput” (actual system output).

Traditional				
Metric Category	Benchmark	AI-Assisted Result	Impact Factor	Citation
Task Completion	Baseline (1.0x)	1.26x - 1.55x Faster	High Positive	(Smit et al., 2024)
Code Review Time	60-90 mins/PR	30-45 mins/PR	High Positive	(Balachandran & Fawzer, 2025)
Context Switch	15-20 mins recovery	Reduced interruption	Medium Positive	(Reddy Vootukuri, 2025)
Debugging Time	High variance	Standardized reduction	High Positive	(Arora, 2025)

Table B1: Aggregated Productivity Metrics from Empirical Studies.

The data indicates a consistent reduction in time-on-task for routine coding activities. Smit et al. (Smit et al., 2024) report significant gains in task completion velocity when developers use tools like GitHub Copilot. Specifically, the “blank page problem”—the initial inertia of starting a new module—is virtually eliminated. Furthermore, Balachandran and

Fawzer (Balachandran & Fawzer, 2025) demonstrate that AI-integrated code review tools significantly reduce the latency of Pull Request (PR) cycles by automating the generation of summaries and initial vulnerability scans.

However, these gains are not uniform. Zuo et al. (Zuo et al., 2024) emphasize that while AI can generate PR titles and descriptions effectively, the *accuracy* of these generations relies heavily on the quality of the diffs and the context provided. If the underlying code changes are complex or poorly structured, the AI's summarization capabilities degrade, potentially requiring more time for human correction than manual writing would have taken.

4.2.2 Security and Supply Chain Vulnerabilities

A critical finding of this thesis is the introduction of new attack vectors through AI-generated code. The data below categorizes the prevalence of specific security risks identified in AI-assisted development environments.

Risk Category	Description	Detection Difficulty	Mitigation Strategy
Adversarial Code	Maliciously prompted injection	High	Enhanced Benchmarks
Hallucination	Non-existent libraries/APIs	Medium	SBOM Verification
Supply Chain	Dependency confusion	High	Blockchain/SBOM
Data Leakage	Training data exposure	Medium	Local LLM Hosting

Table B2: Taxonomy of AI-Introduced Security Risks. Sources: (Swaraj et al., 2025), (Shukla, 2025), (Aideyan et al., 2025).

Swaraj et al. (Swaraj et al., 2025) provide a benchmark dataset revealing that adversarial prompting can trick generic LLMs into generating insecure code patterns that bypass standard static analysis tools. This is particularly dangerous in community-driven platforms like Stack Overflow, where AI-generated answers may propagate vulnerabilities to thousands of developers.

Furthermore, the software supply chain faces new pressures. Shukla (Shukla, 2025) argues that the ease of generating code increases the volume of third-party dependencies included in projects. This necessitates the automated generation and management of Software Bill of Materials (SBOMs). Without automated SBOM management, the opacity of AI-generated codebases makes it nearly impossible to track vulnerability propagation. Aideyan et al. (Aideyan et al., 2025) propose using blockchain-reproducible builds to counter this, ensuring that the provenance of every line of code—whether human or AI-written—is immutable and traceable.

4.2.3 Governance and Compliance Standards

The rapid adoption of GenAI has outpaced regulation, but standards are emerging. The following table outlines the key components of ISO/IEC 42001:2023 as they apply to software engineering organizations.

ISO 42001 Domain	Engineering Application	Compliance Requirement	Citation
Risk Management	AI Code Safety	Auto-testing protocols	(Biroğul et al., 2025)
Data Quality	Training Data Vetting	Clean data pipelines	(Seet, 2025)
Transparency	Explainability	Decision logging	(Biroğul et al., 2025)
Lifecycle Mgmt	Model Updates/Versioning	CI/CD Integration	(Jamili et al., 2025)

Table B3: Application of ISO/IEC 42001:2023 to Software Engineering. Sources: (Seet, 2025), (Biroğul et al., 2025).

Biroğul et al. (Biroğul et al., 2025) emphasize that ISO 42001 provides the first comprehensive framework for managing AI systems organizationally. For software engineering

leaders, this means moving beyond ad-hoc tool adoption to a structured management system that accounts for legal liability and ethical deployment. Seet (Seet, 2025) notes that legal compliance is no longer optional; as AI tools become embedded in critical infrastructure, adherence to these standards will likely become a prerequisite for liability insurance and regulatory approval.

4.3 Appendix C: Glossary of Terms

This glossary defines key technical terms used throughout the thesis, contextualizing them within the specific domain of AI-augmented software engineering.

Adversarial Prompting A technique where malicious inputs are designed to manipulate an AI model into producing harmful, incorrect, or insecure outputs. In the context of software engineering, this involves crafting prompts that cause coding assistants to generate vulnerabilities or bypass security filters (Swaraj et al., 2025).

Agentless Framework A software engineering approach that uses Large Language Models (LLMs) for code generation and repair without the complex state management of autonomous agents. These frameworks typically use a two-phase process (localization and repair) to reduce the cost and complexity associated with fully agentic systems (Xia et al., 2024).

Automated Pull Request (PR) Analysis The use of Generative AI to automatically analyze code changes, generate titles and summaries, and identify potential issues before human review. This technology aims to reduce the cognitive load on maintainers and accelerate the code integration process (Zuo et al., 2024)(Balachandran & Fawzer, 2025).

Context-Aware Code Review An advanced review methodology where the AI tool analyzes not just the syntax of the changed code, but the semantic context of the surrounding codebase, commit history, and project documentation. This allows for more relevant and accurate critiques compared to traditional static analysis (Balachandran & Fawzer, 2025).

Generative AI (GenAI) A class of artificial intelligence systems capable of generating new content (text, code, images) in response to prompts. In software engineering, this primarily refers to Large Language Models (LLMs) trained on vast repositories of source code (e.g., GitHub) to assist in development tasks (Lakshmi et al., 2025)(Esposito et al., 2024).

Human-Centered Software Engineering (HCSE) An approach to software development that prioritizes the cognitive needs, capabilities, and limitations of the human developers and users. In the AI era, HCSE focuses on designing AI assistants that augment rather than replace human decision-making, ensuring that the “human in the loop” maintains agency and understanding (Seffah et al., 2009).

ISO/IEC 42001:2023 An international standard specifying requirements for establishing, implementing, maintaining, and continually improving an Artificial Intelligence Management System (AIMS) within organizations. It provides the governance framework necessary for the safe and compliant adoption of AI tools in enterprise environments (Seet, 2025)(Biroğul et al., 2025).

Large Language Model (LLM) A deep learning algorithm that can recognize, summarize, translate, predict, and generate text and other content based on knowledge gained from massive datasets. Models like GPT-4 and Claude are foundational to tools like GitHub Copilot (Zuo et al., 2024)(Xia et al., 2024).

Software Bill of Materials (SBOM) A formal, machine-readable inventory of software components and dependencies, their hierarchical relationships, and their licensing information. Automated SBOM generation is critical in AI-assisted development to track the provenance of AI-suggested libraries and mitigate supply chain risks (Shukla, 2025)(Syed, 2024).

SWE-bench A rigorous evaluation framework designed to test the capabilities of Language Models on real-world software engineering issues collected from GitHub. It serves

as a standard metric for assessing the ability of AI agents to resolve complex coding tasks autonomously (Zhu & Kang, 2025).

Synthetic Pair Programmer A conceptual metaphor describing the role of AI coding assistants (e.g., GitHub Copilot) as collaborative partners rather than simple tools. This relationship mimics the dynamic of human pair programming, where the AI offers suggestions, completions, and critiques in real-time (Reddy Vootukuri, 2025)(Smit et al., 2024).

4.4 Appendix D: Implementation and Governance Resources

This appendix provides actionable resources for engineering leadership and practitioners regarding the implementation of AI tools. It synthesizes the governance strategies and risk mitigation techniques discussed in the literature review into practical checklists and frameworks.

4.4.1 Strategic Adoption Framework

Implementing GenAI in a software organization requires a structured approach to avoid “shadow AI” usage and ensure security. The following framework, adapted from Barón (Barón, 2025) and Deloitte’s insights (Deloitte, 2024), outlines a four-phase implementation strategy.

Phase	Objective	Key Actions	Success Metric
1. Assessment	Identify high-value use cases	Survey dev teams; Audit current toolchain	Use case clarity
	Test efficacy & security	Deploy to non-critical teams; Sandbox testing	Dev satisfaction

Phase	Objective	Key Actions	Success Metric
3. Governance	Establish policy guardrails	Define acceptable use; Implement ISO 42001	Compliance rate
4. Scale	Broad deployment	Integration with CI/CD; Training programs	Velocity increase

Table D1: Strategic AI Adoption Framework for Engineering Organizations.

The Assessment phase is critical. Organizations must determine where AI adds value versus where it introduces unnecessary risk. For example, applying AI to generate boilerplate code for UI components offers high value with low risk, whereas using AI to generate cryptographic implementation logic carries extreme risk.

4.4.2 Risk Mitigation Checklist

Based on the supply chain security findings by Syed (Syed, 2024) and Aideyan et al. (Aideyan et al., 2025), the following checklist is recommended for all organizations integrating GenAI into their production pipelines.

1. Code Provenance & Supply Chain - [] **Mandatory SBOMs:** All AI-generated code projects must auto-generate a Software Bill of Materials (Shukla, 2025). - [] **Dependency Verification:** Automated scanning of all AI-suggested libraries to prevent “dependency confusion” attacks. - [] **Immutable Builds:** Implementation of blockchain-verified or signed builds to ensure code integrity from commit to deployment (Aideyan et al., 2025).

2. Quality Assurance & Review - [] **Human-in-the-Loop:** Mandatory human review for all AI-generated Pull Requests. No auto-merge for AI code (Cihan et al., 2025). - [] **Context-Aware Scanning:** Utilization of advanced static analysis tools that understand semantic context, not just syntax (Balachandran & Fawzer, 2025). - [] **Adversarial**

Testing: Regular red-teaming of AI assistants using adversarial prompts to check for leaked secrets or insecure patterns (Swaraj et al., 2025).

3. Policy and Compliance - [] Data Privacy Boundaries: Strict prohibition of pasting proprietary logic or PII into public LLM interfaces (Rosenbaum, 2024). - [] **ISO Alignment:** Alignment of internal AI policies with ISO/IEC 42001 standards regarding risk management and transparency (Biroğul et al., 2025). - [] **Training:** Mandatory training for developers on the limitations and hallucination risks of LLMs (Lakshmi et al., 2025).

4.4.3 Future-Readiness: Cloud and Infrastructure

As organizations move toward “Intelligent Cloud Systems,” the infrastructure supporting AI development must evolve. Jamili et al. (Jamili et al., 2025) propose a framework for sustainable and secure AI at scale. This involves “adaptive AI orchestration,” where the underlying cloud infrastructure dynamically allocates resources based on the computational needs of the AI models being used.

For software engineers, this means the development environment itself is becoming “smart.” The IDE is no longer a static text editor but a terminal for an intelligent cloud system that manages context, retrieves relevant documentation via RAG (Retrieval-Augmented Generation), and enforces security policies in real-time. Preparing for this future requires investing in strong cloud architectures that can support the high bandwidth and low latency required for smooth AI interaction.

References

Aideyan, Pesé, & Brooks. (2025). Advancing Automotive Software Supply Chain Security: A Blockchain-Reproducible Build Approach. *SAE technical paper series*. [Https://doi.org/10.4271/2025-01-0456](https://doi.org/10.4271/2025-01-0456).

Ali, & Yue. (2015). Formalizing the ISO/IEC/IEEE 29119 Software Testing Standard. IEEE. (pp. 396-405). [Https://doi.org/10.1109/models.2015.7338271](https://doi.org/10.1109/models.2015.7338271)

Arora. (2025). *How AI Can Help Transform Developer Productivity Through Code Assistants*. Front Matter. [Https://doi.org/10.59350/wxbdd-nfr76](https://doi.org/10.59350/wxbdd-nfr76)

Balachandran, & Fawzer. (2025). Context-aware code review: integrating generative AI for automated pull request analysis. Department of Computer Science & Engineering. [Https://doi.org/10.31705/adscai.2025.53](https://doi.org/10.31705/adscai.2025.53)

Barón. (2025). Towards an Adoption Framework to Foster Trust in AI-Assisted Software Engineering. IEEE. (pp. 247-249). [Https://doi.org/10.1109/cain66642.2025.00038](https://doi.org/10.1109/cain66642.2025.00038)

Biroğul, Şahin, & osgørli. (2025). Exploring the Impact of ISO/IEC 42001:2023 AI Management Standard on Organizational Practices. *Advances in Artificial Intelligence Research*. [Https://doi.org/10.54569/aaair.1709628](https://doi.org/10.54569/aaair.1709628).

Brandebusemeyer. (2025). Interactions with Generative AI: Wearables to Measure Developer Experience and Productivity Objectively. IEEE. (pp. 148-150). [Https://doi.org/10.1109/icse-companion66252.2025.00043](https://doi.org/10.1109/icse-companion66252.2025.00043)

Cihan, Haratian, Içöz, Gül, Devran, Bayendur,... & Tüzün. (2025). Automated Code Review in Practice. *2025 IEEE/ACM 47th International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. [Https://doi.org/10.1109/ICSE-SEIP66354.2025.00043](https://doi.org/10.1109/ICSE-SEIP66354.2025.00043).

Deloitte. (2024). *AI and software development quality / Deloitte Insights*. [Https://www.deloitte.com/us/en/insights/industry/technology/how-can-organizations-develop-quality-software-in-age-of-gen-ai.html](https://www.deloitte.com/us/en/insights/industry/technology/how-can-organizations-develop-quality-software-in-age-of-gen-ai.html)

Esposito, Janes, Taibi, & Lenarduzzi. (2024). Generative AI in Evidence-Based Software Engineering: A White Paper. <Https://doi.org/10.48550/arXiv.2407.17440>

Jamili, Peta, Taware, Krishnan, & Perla. (2025). A Framework for Intelligent Cloud Systems: Enabling Secure, Policy-Driven, and Sustainable AI at Scale. <Https://doi.org/10.1109/I2ITCON65200.2025.11210587>

Lakshmi, Helen, & Sambasivam. (2025). *8 Redefining and Transforming Software Development with Generative AI*. De Gruyter. <Https://doi.org/10.1515/9783111677798-008>

Reddy Vootukuri. (2025). *GitHub Copilot Chat in Developer Workflow*. Apress. Https://doi.org/10.1007/979-8-8688-2196-7_3

Rosenbaum. (2024). *When Medicaid Unwinding Meets AI: In the Matter of DeLoitte Consulting* . Milbank Memorial Fund. <Https://doi.org/10.1599/mqop.2024.0221>

Seet. (2025). *ISO 42001 and the Law*. Apress. Https://doi.org/10.1007/979-8-8688-2099-1_3

Seffah, Vanderdonckt, & Desmarais. (2009). *Human-Centered Software Engineering: Software Engineering Architectures, Patterns, and Sodels for Human Computer Interaction*. Springer London. Https://doi.org/10.1007/978-1-84800-907-3_1

Shukla. (2025). AI-Driven SBOM: Automated Software Bill of Materials Generation and Management. *Frontiers in Emerging Artificial Intelligence and Machine Learning*, 02(12), 86-92. <Https://doi.org/10.64917/feaiml/volume02issue12-08>.

Smit, Smuts, Louw, Pielmeier, & Eidelloth. (2024). The impact of GitHub Copilot on developer productivity from a software engineering body of knowledge perspective. *Americas Conference on Information Systems*. <Https://www.semanticscholar.org/paper/d42ed56a876c314bd4495942a4468c49720c68d0>.

Swaraj, Agarwal, Joshi, & Kumar. (2025). Detecting Adversarial Prompted AI-Generated Code on Stack Overflow: A Benchmark Dataset and an Enhanced Detection Approach. *IEEE International Conference on Software Maintenance and Evolution*. <Https://doi.org/10.1109/ICSME64153.2025.00089>.

Syed. (2024). *Emerging Trends in Software Supply Chain Security*. Apress. Https://doi.org/10.1007/979-8-8688-0799-2_10

Ulfnes, Moe, Stray, & Skarpen. (2024). *Transforming Software Development with Generative AI: Empirical Insights on Collaboration and Workflow*. Springer Nature Switzerland. Https://doi.org/10.1007/978-3-031-55642-5_10

Wang. (2025). Sensory Experience-Driven CMF Design for Smart Cabins: A Generative AI Collaborative Approach with the MINI Aceman Ocean Wave Green as a Case Study. *The Korean Society of Science & Art*. <Https://doi.org/10.17548/ksaf.2025.06.30.347>.

Xia, Deng, Dunn, & Zhang. (2024). Agentless: Demystifying LLM-based Software Engineering Agents. <Https://doi.org/10.48550/arXiv.2407.01489>

Zhu, & Kang. (2025). UTBoost: Rigorous Evaluation of Coding Agents on SWE-Bench. Association for Computational Linguistics. (pp. 3762-3774). <Https://doi.org/10.18653/v1/2025.acl-long.189>

Zuo, Lan, & Liao. (2024). Exploring the Potential of Large Language Models in Automatic Pull Request Title Generation: An Empirical Study. IEEE. (pp. 191-200). <Https://doi.org/10.1109/apsec65559.2024.00030>