

The Illusion of Competence: Defining “Epistemic Debt” in the Era of LLM-Assisted Software Engineering

Ludovic A. NGABANG

Independent Researcher
aggax27@gmail.com

January 2, 2026

Abstract

Abstract. The integration of Large Language Models (LLMs) into the software development lifecycle represents a shift from *constructive* programming to *curated* programming. While current metrics focus on productivity gains and syntactical correctness, this paper argues that these metrics are insufficient to capture the long-term systemic risks introduced by AI. We propose the concept of **Epistemic Debt**: the divergence between the complexity of a software system and the developer’s cognitive model of that system. Unlike traditional Technical Debt, which is often a conscious trade-off, Epistemic Debt is an invisible accumulation of “unearned” code that functions correctly but lacks a human owner who understands its causality. This paper provides a theoretical framework for this phenomenon, classifies the specific architectural erosions caused by stochastic code generation, and proposes a “Cognitive Ratchet” methodology to mitigate the collapse of maintainability.

1 Introduction

In his seminal 1986 essay, Fred Brooks argued that there is “No Silver Bullet” in software engineering because the essential difficulty of software lies in its complexity, conformity, changeability, and invisibility [1]. For decades, the bottleneck in software production was the translation of human intent into formal syntax.

Today, Large Language Models (LLMs) such as GPT-4 and Claude 3 have effectively removed this syntax barrier. However, by reducing the friction of code generation, we have inadvertently removed the primary mechanism by which developers build mental models: the act of writing itself.

We are entering an era of *Illusory Fluency*. Developers can now produce systems that exceed their own understanding. This paper posits that the primary risk of AI in Software Engineering (AI4SE) is not the introduction of bugs, but the introduction of plausible, working code that

is *epistemically opaque*. We define this accumulation of un-understood logic as **Epistemic Debt**.

2 Theoretical Framework: From Technical to Epistemic Debt

Ward Cunningham’s original definition of Technical Debt (1992) referred to a strategic decision: shipping imperfect code now to speed up delivery, with the intention of refactoring later [2]. It implies agency and awareness.

2.1 The Loss of Agency

In AI-assisted development, the developer often acts as a reviewer rather than an author. When an LLM generates a 50-line algorithm, and the developer accepts it because “it passes the unit tests,” the code enters the codebase without the developer necessarily understanding the edge cases or the algorithmic constraints.

2.2 Defining Epistemic Debt (D_e)

We formalize Epistemic Debt (D_e) as the integral of the difference between the System Complexity (C_s) and the Cognitive Grasp (G_c) of the maintenance team over time:

$$D_e = \int_{t=0}^T (C_s(t) - G_c(t)) dt$$

In traditional development, C_s and G_c rise in tandem. In AI-assisted development, C_s can rise exponentially while G_c remains flat. When D_e exceeds a critical threshold, the system becomes “Legacy Code” the moment it is written—unmaintainable and terrifying to touch.

3 Taxonomy of AI-Induced Architectural Erosion

We categorize the manifestations of Epistemic Debt into three distinct phenomena.

3.1 1. The Stochastic Spaghetti Effect

Human architects tend to think in patterns and modular structures (Design Patterns). LLMs, being probabilistic token predictors, optimize for local probability rather than global coherence.

- **Symptom:** The model solves Problem A and Problem B in the same file using two completely different coding styles or abstraction levels, simply because the prompts were different.
- **Result:** A breakdown of "Conceptual Integrity," which Brooks identified as the most important consideration in system design.

3.2 2. The Mirage of Correctness (Dependency Hallucination)

LLMs often invent APIs that *should* exist but do not. This forces the developer to write "shim" code or wrappers to make reality match the hallucination. *Deep Analysis:* This reverses the dependency flow. Instead of the tool serving the architecture, the architecture is warped to accommodate the stochastic output of the tool.

3.3 3. Context Window Amnesia

Even with large context windows (128k+ tokens), LLMs struggle with the "unknown unknowns" of a large codebase. They cannot reason about the side effects a generated function might have on a distant module not included in the prompt context. *Result:* The introduction of "Heisenbugs"—bugs that disappear or alter behavior when one attempts to probe them, often caused by subtle state mismanagement generated by the AI.

4 Proposed Methodology: The Cognitive Ratchet

To survive in an AI-saturated ecosystem, Software Engineering must pivot from *Generation* to *Validation*. We propose a mechanism called the "Cognitive Ratchet."

4.1 Principle of Inverse Documentation

Currently, developers ask AI to write code. To reduce Epistemic Debt, the flow must be reversed:

"Do not commit AI code until you can explain it to the AI."

Developers should force the LLM to generate code, and then the developer must write the documentation *manually*. If the developer cannot document the logic clearly, the code is rejected. This forces the cognitive synchronization required to lower D_e .

4.2 Intent-Based Testing

Unit tests usually verify *implementation*. In the AI era, we need "Property-Based Testing" (e.g., Hypothesis in Python) that verifies *invariants*. Since we cannot trust the implementation path chosen by the AI, we must rigorously constrain the boundaries of its behavior.

5 Conclusion

The "Genius" of the future software engineer will not lie in their ability to implement quicksort or invert a binary tree—tasks now commoditized by AI. Instead, the value shifts to **Systemic Curation**.

If we ignore Epistemic Debt, we risk creating a global infrastructure of software that functions by statistical probability rather than deterministic logic. We are building a Tower of Babel where the bricks are laid by machines, and we, the architects, no longer speak the language of construction. The challenge of the next decade is not to generate more code, but to understand what we have already generated.

References

- [1] F. P. Brooks, "No Silver Bullet: Essence and Accidents of Software Engineering," *Computer*, vol. 20, no. 4, pp. 10–19, 1987.
- [2] W. Cunningham, "The WyCash portfolio management system," in *OOPSLA '92 Addendum*, 1992.
- [3] D. L. Parnas, "On the criteria to be used in decomposing systems into modules," *Communications of the ACM*, vol. 15, no. 12, 1972.
- [4] A. Vaswani et al., "Attention Is All You Need," *Advances in Neural Information Processing Systems*, 2017.
- [5] N. Perry, D. Srivastava, D. Kumar, and D. Boneh, "Do Users Write More Insecure Code with AI Assistants?," *ACM CCS*, 2023.