

Peer Review

Review of: "The LLM Productivity Cliff: Threshold Productivity and AI-Native Inequality"

Tyler Thomas Procko¹

1. Embry-Riddle Aeronautical University, United States

This is a well-written conceptual piece that serves to articulate, in very clear language, an important pattern, that is, thresholds in LLM productivity. It is important because, in the current zeitgeist, it is all too common to assume that LLM use guarantees a productivity boost.

The article does a solid job of connecting and synthesizing disparate findings across domains into a coherent narrative. The most important contribution is that the defined engineering mindset of “architectural literacy” moves us beyond the “better prompting” advice so common to most of these discussions.

That being said, the article is more of a hypothesis-generating commentary than a rigorous empirical piece, which is perfectly acceptable because it is framed as such (see the intro of section 5). The reason I assert this is that the evidence base is thin (a handful of 2025 studies, some of which are not peer-reviewed, and an Anthropic podcast citation), and there is no longitudinal study (but that is noted in section 5 as well, and will surely be remedied as time proceeds). Additionally, one possible issue is that the piece may overgeneralize the threshold effects as if they apply universally; but the cited evidence shows that they are, realistically, domain- and task-dependent. Customer support, for instance, shows different patterns than coding, but they seem to be treated as variations of a single phenomenon.

Overall, a wonderful read and very insightful for anyone engaged in LLM use as a software engineer. Surely there is wisdom to be gained from reading this, and the incentive to use LLMs to their fullest potential atop scaffolding/tooling, rather than as simple chatbots.

Declarations

Potential competing interests: No potential competing interests to declare.