



Methodological Approaches to Assessing Productivity and Resource Planning in Hybrid Human-AI Software Engineering Teams

Andrii Shaliev

Senior Delivery Director, Client Partnership and Growth @ Trinetix Inc., Nashville, TN, USA.

Abstract

The article is dedicated to the analysis of methodological approaches to assessing productivity and resource planning in hybrid human-AI software engineering teams. The relevance of the study is determined by the rapid integration of generative AI tools into development processes and the growing mismatch between traditional productivity metrics and emerging socio-technical realities. Scientific novelty lies in the analytical reinterpretation of productivity as a dynamic balance between generative speed, stabilization effort, architectural coherence, and shared understanding, rather than as a single output indicator. The work describes how AI reshapes coordination rhythms, redistributes cognitive and organizational effort, and alters the temporal structure of planning and risk manifestation. Special attention is paid to the delayed effects of AI-driven acceleration, including architectural drift, erosion of contextual knowledge, and the growing invisibility of preventive work. The study sets itself the goal of identifying methodological principles suitable for evaluating productivity and planning capacity in hybrid teams. To achieve this goal, analytical synthesis, comparative analysis, and source-based conceptual modeling are applied. The conclusion outlines implications for measurement frameworks and governance practices. The article will be useful for researchers, software engineering managers, and practitioners involved in AI-augmented development environments.

Keywords: Hybrid Human-AI Teams, Software Engineering Productivity, Resource Planning, Generative AI, Stabilization Work, Architectural Governance, Coordination Mechanisms, Productivity Metrics.

INTRODUCTION

The growing adoption of generative AI in software engineering has intensified long-standing debates about how productivity should be measured and planned in complex development environments. While AI tools demonstrably accelerate code generation, their integration exposes structural tensions between speed, quality, coordination, and long-term system stability. Classical productivity models, grounded in linear relationships between effort and output, increasingly fail to capture these dynamics [1].

The purpose of this article is to develop a methodological perspective on productivity assessment and resource planning in hybrid human-AI software engineering teams. The study addresses the following objectives:

1) to analyze how AI-driven acceleration reconfigures the distribution of effort across development phases;

- 2) to identify methodological limitations of traditional productivity and planning metrics in hybrid teams;
- 3) to systematize analytical directions suitable for evaluating productivity under conditions of probabilistic automation.

The novelty of the study lies in treating productivity not as a static performance indicator, but as a moving boundary shaped by socio-technical coordination, delayed risk effects, and evolving forms of accountability and learning within hybrid teams.

METHODS AND MATERIALS

The materials for this study consist of recent scholarly works devoted to generative AI in software engineering, human-AI collaboration, productivity measurement, organizational readiness, and hybrid intelligence. The study by Ankush Sharma [1] examines productivity shifts and technical debt accumulation

Citation: Andrii Shaliev, "Methodological Approaches to Assessing Productivity and Resource Planning in Hybrid Human-AI Software Engineering Teams", Universal Library of Engineering Technology, 2026; 3(1): 54-59. DOI: <https://doi.org/10.70315/uloap.ulete.2026.0301009>.

in AI-driven engineering teams. The work of Viktoria Stray et al. [2] analyzes coordination risks and failure patterns associated with AI-assisted development. Leonardo Banh et al. [3] explore structural transformations of software engineering practices under generative AI adoption. Silvia Abrahão et al. [4] conceptualize human-centered software engineering in the AI era. Erik Brynjolfsson et al. [5] investigate the productivity effects of generative AI across knowledge-intensive work. The study by Fan et al. [6] focuses on learning and collaboration outcomes in AI-assisted pair programming. Sauer and Burggräf [7] propose a framework for hybrid intelligence alignment. Aldoseri et al. [8] assess organizational readiness for AI-based transformation, while Madanchian et al. [9] analyze AI-based human resource management tools.

To write the article, analytical synthesis, comparative analysis, conceptual modeling, and qualitative source analysis were used. These methods enabled the integration of heterogeneous empirical and conceptual findings into a unified methodological framework.

RESULTS

The configuration of productivity inside hybrid human-AI teams reveals itself not as acceleration, but as a re-shaping of where effort accumulates and where it disappears. Early signals come from the micro-timing of everyday work: developers spend less time writing boilerplate and more time repairing consequences that only become visible after integration. A team may move faster on day one and feel slower on day forty. This temporal dislocation unsettles classical planning models.

In mixed teams, velocity fractures into two distinct rhythms that rarely align. Generative velocity grows when tasks are modular, well specified, and textually legible; integration velocity decelerates when outputs must be reconciled with legacy architectures, data contracts, and security boundaries. In feature development, several teams recorded 40% gains in visible output, while debugging, integration, and maintenance of AI-generated code slowed by 25% during the same sprints [1]. Planning boards still counted story points, yet the meaning of a “completed story” drifted: completion at merge time no longer matched completion at production time. The planning horizon, therefore, oscillated between short bursts of speed and delayed repair cycles that materialized two to three sprints later, when technical debt crystallized into incidents.

Beneath this oscillation lies a deeper methodological problem: productivity metrics travel along different axes than risk metrics. When AI tool usage rose across delivery pipelines, delivery stability declined, primarily through architectural drift rather than coding errors [2]. The data topology becomes inverted - local efficiency improvements coincide with system-level fragility. The systematization of quantitative effects is presented below (Figure 1).

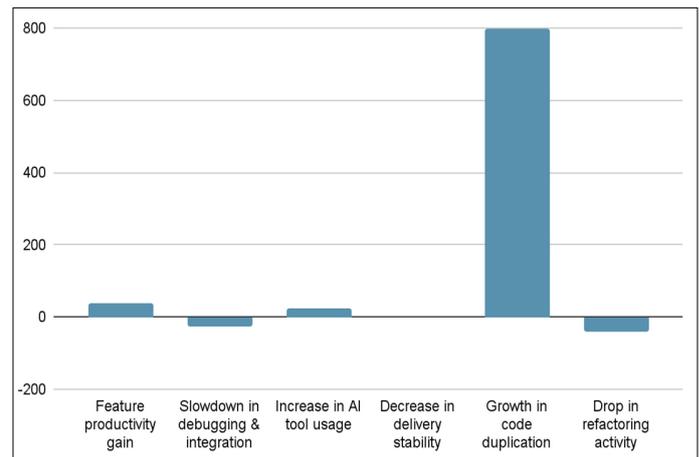


Figure 1. Quantitative effects of AI adoption on productivity and stability (compiled by the author based on [1,2])

Teams that optimized for rapid pull requests often created hidden coupling between services that only surfaced under load or during deployments. Speed looked rational at the micro level and irrational at the macro level. No dashboard reconciled these scales.

A second trajectory centers on coordination rather than speed. AI tools flatten expertise gradients by giving junior developers immediate access to plausible solutions, yet they simultaneously erode shared architectural memory. A 2025 survey of 18 major CTOs found that 16 had already experienced production disasters triggered by unvetted AI-generated code [2]. The crisis was rarely about syntax; it stemmed from missing context - implicit design rules that had never been written down. When lead engineers left, teams discovered that nobody could reconstruct why certain boundaries existed. Knowledge became episodic, stored in prompts rather than in collective reasoning.

This drift becomes measurable in code repositories. Analysis of more than 200 million lines of code shows an 8× increase in duplicated blocks after mass AI adoption, accompanied by a predicted global technical debt of \$1.5 trillion by 2027 [1]. In the same datasets, refactoring activity dropped by 40%, while raw duplication jumped relative to pre-AI baselines [1]. The pattern suggests that AI optimizes for local adequacy, not structural coherence. Systems grow wider rather than cleaner. Short sentences accumulate like brittle scaffolding. The architecture ages without maturing.

Teams respond by inventing compensatory governance rather than abandoning AI. Some introduce a ground-rules file at the project root that explicitly states prohibited patterns and mandatory conventions. Others maintain module-specific markdown guides so that prompts carry architectural memory. Automated validators begin to police boundaries, flagging when AI crosses predefined service lines. These practices do not restore the past; they formalize the friction between speed and stability. A rule set replaces intuition, and the cost of coordination migrates from informal conversations to toolchains.

A third trajectory concerns capacity planning. Traditional models assume linear relationships between effort and output; hybrid teams break that assumption. AI proficiency varies sharply across individuals, tasks, and tools, producing uneven micro-productivity that destabilizes resource allocation. Teams start to decompose velocity into generation and stabilization components, treating them as separate

planning currencies. Capacity calibration then compares story-point throughput with the percentage of commits created with AI assistance, code-review cycle time, bug-escape rates, and pull-request merge frequency. Planning becomes multi-metric by necessity, not by theory. The systematization of methodological directions is presented below (Table 1).

Table 1. Typology of methodological directions for assessing productivity in hybrid human-AI teams (compiled by the author based on [1-9])

Methodological direction	Core analytical focus
Velocity decomposition	Separation of code generation work and stabilization work into distinct planning dimensions rather than a single measure of throughput
Socio-technical coordination	Interaction between human roles, AI tools, repositories, and automated validators as a unified productive system
Architectural governance	Mechanisms that constrain architectural drift, service boundary violations, and uncontrolled code duplication
Learning and skill formation	How AI reshapes developers' cognitive habits, reflective practice, and depth of understanding
Accountability and ownership	Shift from formal authorship toward epistemic responsibility and system comprehension
Risk visibility and delay effects	Temporal lag between early productivity gains and later system failures or incidents
Hybrid intelligence alignment	Continuous calibration between human judgment and AI automation in decision-making
Organizational readiness for AI	Institutional capacity to absorb AI-driven change in planning, governance, and evaluation

Yet numbers alone mislead. When AI accelerates initial coding, the psychological posture of developers shifts from authorship to verification. Review labor expands in ways that are hard to quantify. One additional hour spent reviewing ambiguous AI output may prevent weeks of rework, but the benefit appears only in the counterfactual - incidents that never happen. Planning frameworks struggle with invisible savings. The spreadsheet remains flat while the system becomes safer.

A fourth trajectory highlights learning and skill formation. In educational settings, AI-assisted pair programming raises student motivation and reduces programming anxiety compared with traditional pair programming or solitary work, while also improving collaborative performance [6]. The effect is not simply about speed; it reconfigures how novices engage with uncertainty. They iterate more, explain less, and treat code as negotiable rather than authoritative. Over time, this can both democratize access to complex tasks and weaken deep conceptual understanding.

At the professional level, a similar tension appears. Developers gain rapid access to explanations and examples, yet risk losing the habit of slow reasoning. Some teams report a subtle decline in architectural literacy as AI handles more surface complexity. Others counteract this by requiring architects to write explicit Architectural Decision Records for every significant pull request, making reasoning visible again. Learning migrates from tacit apprenticeship to documented trace.

A fifth trajectory reframes accountability. When AI generates most of a pull request, authorship blurs into stewardship.

Responsibility shifts from “who wrote the code” to “who understands the system well enough to justify it.” Teams experiment with tiered ownership models, disclosure of AI usage in commits, and reviewer contracts that bind explanation to approval. Incident rules begin to prioritize depth of understanding over formal approval chains. Ownership becomes epistemic rather than procedural.

Across all trajectories, the boundary between productivity and risk remains unstable. Generative systems can handle non-routine tasks that once required scarce expertise, often producing coherent solutions to complex problems with minimal prompting [5]. At the same time, probabilistic behavior introduces unpredictable failure modes that amplify coordination costs when stakes are high. The hybrid team becomes a site of continuous calibration rather than settled efficiency.

Several cross-cutting patterns converge. First, task suitability matters: well-scoped, low-ambiguity tasks absorb AI gains, while open-ended design work amplifies downstream repair costs. Second, context depth matters: tools integrated tightly with repositories perform better than standalone chat interfaces, yet raise data-protection concerns. Third, timing matters: benefits arrive early, liabilities arrive later. Fourth, governance matters: lightweight rules and automated checks reduce drift without eliminating it.

Seen together, productivity in hybrid teams no longer equates to faster code creation. It becomes a balance among generation speed, stabilization effort, architectural coherence, and shared understanding. The most capable teams do not chase maximum velocity; they cultivate predictable friction. They

slow down in the right places, formalize invisible knowledge, and treat AI as a catalyst that exposes hidden dependencies rather than a replacement for them.

A final tension persists. When AI accelerates routine work, organizations can redistribute capacity toward higher-value problems, yet they may also compress schedules and expectations, converting gains into pressure rather than slack. Productivity improvements then coexist with rising stress and narrower margins for error. Stability emerges not from technology alone, but from how teams reorganize around it.

In this sense, the methodological challenge is less about measuring output and more about mapping where effort migrates. Hybrid productivity is a moving boundary between speed, stability, and shared comprehension. The systematization of interactions is presented below (Figure 2).

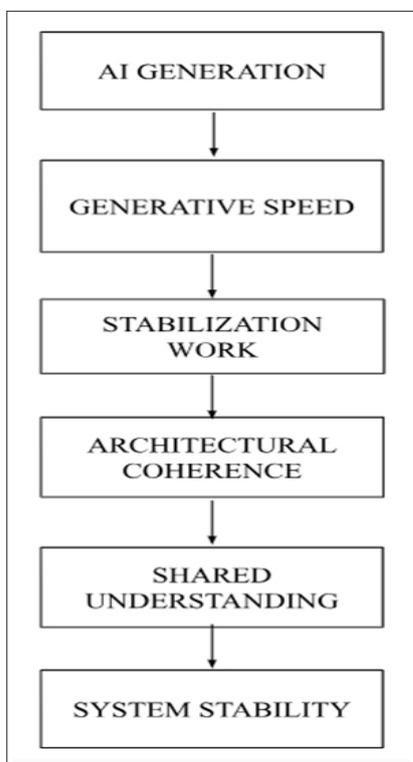


Figure 2. Conceptual scheme of productivity dynamics in hybrid human-AI teams (compiled by the author based on [1,2,3,7])

It resists single indicators and rewards composite, situated measurement frameworks that track generative velocity, stabilization work, architectural drift, and learning simultaneously.

DISCUSSION

The methodological frame for assessing productivity in hybrid human-AI software teams begins to wobble the moment measurement treats output as a single, stable quantity. The analytic problem is not that existing metrics are wrong; it is that they presuppose a unitary locus of value. In practice, effort migrates across phases, roles, and artifacts

that traditional planning instruments cannot see. A sprint burndown chart may close cleanly while the architecture grows less legible. The tension between local efficiency and systemic intelligibility becomes the central interpretive friction.

The resource-planning apparatus inherited from classical agile practices assumes that variability is noise around a predictable mean. In mixed human-AI settings, variability turns into the signal itself. Fluctuations in code quality, review load, and stabilization effort are not random disturbances but structural effects of generative tools. The planning horizon stretches backward and forward at once: backward because hidden dependencies resurface through incidents, forward because teams must anticipate forms of debt that are not yet visible. Midway through a release cycle, the very definition of “done” starts to shift, and planners are forced to renegotiate completion criteria while execution is already underway.

At a deeper level, the coordination rhythm inside teams changes shape. Organizations initially treat AI as a productivity amplifier that preserves existing divisions of labor. Gradually, the boundary between creating, reviewing, and integrating dissolves. Reviewers become co-architects of AI outputs, developers turn into curators of prompts, and managers begin to arbitrate conflicts between speed and coherence rather than allocating hours. Planning tools calibrated for role specialization struggle with this hybridization. It becomes visible that capacity is no longer primarily a function of person-days but of alignment among humans, models, repositories, and automated validators.

Methodological friction also arises from how risk is rendered legible. Traditional dashboards privilege proximal indicators - cycle time, pull-request throughput, deployment frequency - because they are easy to collect. Hybrid systems displace risk into less observable zones: architectural drift, semantic misalignment, and erosion of collective memory. The evidentiary signal appears only after a delay, often triggered by leadership transitions, staff turnover, or scale events. Risk assessment must therefore move from static snapshots to longitudinal pattern detection, yet most organizations lack instruments that track knowledge continuity or contextual depth.

Another boundary emerges around learning. Educational and professional settings reveal different effects of AI assistance, yet both unsettle assumptions about skill acquisition. In classrooms, generative tools lower affective barriers and accelerate experimentation; in industry, they can compress reflective practice. The analytical frame for productivity measurement rarely distinguishes between speed of task completion and durability of understanding. A team that ships faster today may cultivate a thinner substrate of expertise tomorrow. Planning models that ignore this temporal trade-off conflate immediate throughput with long-term capability.

Governance mechanisms illustrate a further methodological dilemma. Ground-rules files, architectural validators, and modular pull-request limits are often introduced as technical fixes. In reality, they redistribute power and responsibility. Decision authority migrates toward those who design rules and tooling rather than those who write code. The institutional arrangement begins to privilege process architects over domain experts. Productivity metrics calibrated to coding speed miss this shift entirely, yet it shapes outcomes more decisively than any isolated performance number.

Measurement practices also confront a paradox of visibility. Stabilization work - debugging, integration, refactoring, documentation - becomes both more critical and less legible as AI generates larger volumes of plausible code. Time spent preventing incidents rarely registers as achievement. Teams are incentivized to underinvest in invisible labor because dashboards reward what is countable. A methodological reframing is required in which preventive work is treated as productive output rather than slack or overhead.

Midway through this analysis, a different tension surfaces. Some organizations respond by multiplying metrics: separate indicators for generative output, review burden, architectural compliance, and incident exposure. The dashboard grows denser, yet coherence does not follow. More numbers do not automatically yield better interpretation. The problem shifts from data scarcity to analytic overload, where managers drown in signals that pull in incompatible directions.

Alternative framings suggest that productivity should be conceptualized as a dynamic equilibrium among three domains: generative speed, structural integrity, and shared comprehension. Each domain exerts countervailing pressure on the others. Pushing speed without reinforcing integrity produces drift; prioritizing integrity without cultivating comprehension breeds bureaucratic rigidity; centering comprehension without enabling generative tools risks stagnation. No single metric can stabilize this triangle.

Leadership transitions expose these dynamics with unusual clarity. New leaders often misread surface stability as systemic health and reallocate resources toward visible throughput, amplifying latent fragilities. The coordination rhythm becomes brittle long before incidents materialize. Methodologically, this indicates that assessment frameworks must incorporate temporal lags and relational indicators - trust trajectories, decision velocity, and continuity of informal knowledge - alongside technical measures.

A persistent blind spot concerns the unit of analysis. Many studies and dashboards oscillate between individual productivity and team outcomes without clarifying the linkage. In hybrid settings, individual gains do not aggregate linearly; they recombine through codebases, dependencies, and review networks. The analytical frame must treat the repository itself as an active participant in productivity, not merely a container for artifacts.

Finally, an unresolved friction remains around accountability. As authorship blurs into stewardship, responsibility becomes epistemic rather than procedural. Measurement systems built around commit histories or approval chains cannot capture who truly understands the system. Yet replacing them with qualitative assessments of understanding introduces subjectivity that managers resist. The verification regime sits uneasily between quantification and judgment.

These tensions do not converge into a neat synthesis. Instead, they point toward a plural methodology that blends longitudinal tracing, artifact analysis, and organizational diagnostics. Productivity in hybrid human-AI teams appears less as a quantity to be maximized and more as a fragile alignment that must be continually negotiated. The analytic task is not to resolve this instability but to make its contours visible enough for deliberate action.

CONCLUSION

The conducted analysis demonstrates that productivity in hybrid human-AI software engineering teams cannot be adequately assessed using single-dimensional or purely output-oriented metrics. The first objective was fulfilled by showing that AI redistributes effort from code generation toward stabilization, review, and governance activities. The second objective was addressed through the identification of temporal delays, architectural drift, and invisible preventive work that undermine traditional planning models. The third objective was achieved by systematizing methodological directions that account for coordination, learning, accountability, and hybrid intelligence alignment.

The findings indicate that effective resource planning in hybrid teams requires composite measurement frameworks capable of capturing both immediate gains and delayed systemic effects. Productivity emerges as a negotiated equilibrium rather than a maximized quantity, demanding continuous recalibration of metrics, governance mechanisms, and organizational expectations. The study contributes to the emerging theory of hybrid software engineering by reframing productivity as a socio-technical equilibrium rather than a performance output.

REFERENCES

1. Sharma, A. (2023). Developer productivity in AI-driven engineering teams. *World Journal of Advanced Research and Reviews*, 18(2). <https://doi.org/10.30574/wjarr.2023.18.2.0898>
2. Stray, V., Barbala, A., & Wivestad, V. T. (2025). Human-AI collaboration in software development: A mixed-methods study of developers' use of GitHub Copilot and ChatGPT. In *Companion Proceedings of the 33rd ACM International Conference on the Foundations of Software Engineering (FSE Companion '25)* (pp. 1-8). ACM. <https://doi.org/10.1145/3696630.3730566>

3. Banh, L., Holldack, F., & Strobel, G. (2025). Copiloting the future: How generative AI transforms software engineering. *Information and Software Technology*, 183, 107751. <https://doi.org/10.1016/j.infsof.2025.107751>
4. Abrahão, S., Grundy, J., Pezzè, M., Storey, M.-A., & Tamburri, D. A. (2025). Software engineering by and for humans in an AI era. *ACM Transactions on Software Engineering and Methodology*, 34(5). <https://doi.org/10.1145/3715111>
5. Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
6. Fan, G., Liu, D., Zhang, R., Yang, Y., & Liang, L. (2025). The impact of AI-assisted pair programming on student motivation, programming anxiety, collaborative learning, and programming performance: A comparative study with traditional pair programming and individual approaches. *International Journal of STEM Education*, 12, 16. <https://doi.org/10.1186/s40594-025-00537-3>
7. Sauer, C. R., & Burggräf, P. (2025). Hybrid intelligence – Systematic approach and framework to determine the level of human-AI collaboration for production management use cases. *Production Engineering*, 19, 525–541. <https://doi.org/10.1007/s11740-024-01326-7>
8. Aldoseri, A., Al-Khalifa, K. N., & Hamouda, A. M. (2024). Methodological approach to assessing the current state of organizations for AI-based digital transformation. *Applied System Innovation*, 7(1), 14. <https://doi.org/10.3390/asi7010014>
9. Madanchian, M., Taherdoost, H., & Mohamed, N. (2023). AI-based human resource management tools and techniques: A systematic literature review. *Procedia Computer Science*, 229, 367–377. <https://doi.org/10.1016/j.procs.2023.12.039>