

# Human Centric AI Agents for Software Development: A Review of Emerging Paradigms and Empirical Efficacy

HAMD BIN FAYYAZ BAIG, FAYYAZ AHMAD

*Journal of Secondary and Undergraduate Research*

*Published on May 12<sup>th</sup>, 2026*

---

This review article synthesizes ten pivotal publications from 2025 and 2026 to evaluate the evolution of human-centric Artificial Intelligence agents within software engineering. It identifies a transition from reactive auto-complete tools to proactive Multi-Agent Systems and explores the “orchestration” paradigm shift where developers manage autonomous agents. Central to this evaluation is the tension between productivity gains highlighted by the Human-in-the-Loop LLM-based Agent framework’s 82% plan approval rate and the cognitive disruption of the developer’s flow-state. By analyzing the Decentralized Planning Decentralized Execution mechanism, the study establishes that while proactive agents such as “CodeAct” (achieving a 79.3% success rate on the HumanEvalFix benchmark) offer significant efficiency, their real-world integration is hampered by “cognitive lag.” The review further addresses technical deskilling in junior developers, contrasting the high-level intent of “vibe coding” with the need for underlying logical awareness. It argues for AI-based verification and Social Transparency to maintain trust in semi-automated environments. Strategic recommendations include architectural patterns such as context awareness and progressive complexity disclosure to mitigate workflow friction. Furthermore, the evaluation critiques the gap between single-file research prototypes and enterprise-scale multi-file interdependencies, advocating for sandboxed execution environments such as OpenHands. The paper concludes that agentic systems expand rather than replace human capabilities, provided that multi-modal verification frameworks and social transparency protocols are prioritized to ensure accountability. This synthesis provides a roadmap for achieving Hybrid Intelligence through standards of reliability and human-agent synergy.

---

## 1 Introduction

2 In the current technological landscape, Artificial Intelligence (AI) has evolved far beyond the rudimentary predictive capabilities of IntelliSense. We are witnessing a systemic evolution from basic code-completion tools to proactive agents capable of “agentic reasoning.” This signifies a system’s ability to function as an autonomous partner, capable of prospective planning, reasoning through architectural hurdles, and managing complex life cycles with independent initiative. This evolution marks the transition toward “Hybrid Intelligence,” a model defined by the synergy of human strategic creativity and algorithmic pattern recognition, rather than the wholesale replacement of human labour.

15 For a comprehensive evaluation, it is necessary to establish technical nomenclature. Large Language Models (LLM) represent complex architectures trained on vast datasets to facilitate reasoning and multi-step task execution. Multi-Agent Systems (MAS) refer to collectives of specialized AI agents organized into functional units to facilitate a division of labour analogous to a human engineering team. This paradigm shift was recently evaluated by X. Wang *et al.* (2025), who defined proactive agents as systems capable of internalizing project context to anticipate developer requirements. Subsequently, S. Otoum *et al.* (2026) established that this agent-centric methodology provides superior quality control through automated testing frameworks. The objective of this document is to evaluate

whether the substantial productivity benefits of agentic automation outweigh the associated challenges of technical deskilling and cognitive lag. This review examines the efficiency of integrating these agents into professional workflows while accounting for the significant cognitive costs of focus disruption.

The objective of this document is to evaluate whether the substantial productivity benefits of agentic automation outweigh the associated challenges of technical deskilling and cognitive lag. This review examines the efficiency of integrating these agents into professional workflows while accounting for the significant cognitive costs of focus disruption.

## 2 Impacts Of Proactive AI Agents On Developer Productivity

Empirical evidence from 2025–2026 frameworks indicates that proactive AI support offers a non-trivial improvement in developer efficiency compared to traditional prompt-driven paradigms. A primary exemplar is the Human-in-the-Loop LLM-based Agent (HULA) framework, which employs a Decentralized Planning Decentralized Execution (DPDE) mechanism [7]. DPDE bifurcates the cognitive load by separating high-level task planning from low-level code execution across specialized units, facilitating higher precision in complex environments.

When deployed within enterprise-grade platforms such as Atlassian JIRA, HULA achieved an 82% plan approval rate and a 59% pull request merge rate [7]. This suggests that agentic systems can successfully navigate the rigorous gatekeeping and human approval processes inherent in production-ready environments. Furthermore, the "CodeAct" agent, integrated within the OpenHands platform, demonstrated a 79.3% success rate on the HumanEvalFix benchmark, which specifically evaluates an AI's capacity for bug remediation [10]. These agents interact with the development environment through Command-Line Interfaces (CLI) and web browsers within "sandboxed" environments and isolated digital spaces that permit code execution without jeopardizing the primary system integrity.




Figure 1: Comparison of HULA framework performance in JIRA deployments.

## 3 Barriers To Successful Multi-Agent Implementation

Despite empirical gains, the adoption of MAS introduces psychological and technical friction. The "Codellaborator" study conducted by J. Pu *et al.* (2025) utilized a within-subject design involving 18 participants to analyze the trade-offs of proactive assistance. The researchers identified that while proactivity enhances speed, it frequently disturbs the developer's "flow state" the psychological condition of deep immersion.

Such interruptions result in "cognitive lag," a switching cost where the developer must pause their mental process to evaluate the agent's input before attempting to resume the task. If this lag is excessive, the temporal savings provided by the AI are negated. To mitigate these barriers, current research advocates for "presence indicators" visual cues such as progress bars and "interaction scopes," which allow developers to establish digital boundaries where agentic interjection is prohibited. Evidence suggests these features permit developers to incorporate suggestions only at natural cognitive break-points, thereby reducing workflow fragmentation [5].

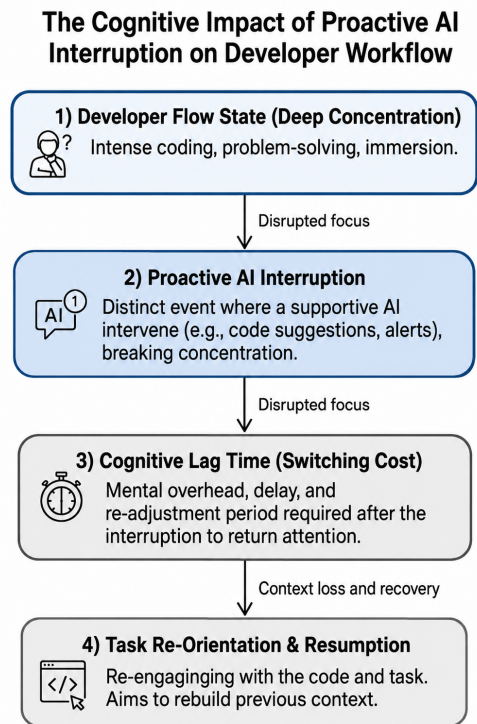


Figure 2: Figure caption goes here.

## 4 The Crisis Of Technical Deskilling In Junior Developers

The transition to higher abstraction layers introduces a risk of technical deskilling, particularly among junior cohorts. L. Chen *et al.* (2025) distinguished between "vibe coding" wherein developers define high-level intent and the rigorous understanding of underlying architecture. While vibe coding facilitates rapid prototyping, the delegation of implementation logic to agents can result in a loss of "situational awareness." If technical details are obscured "under the hood," developers may become incapable of debugging systemic failures or optimizing performance-critical code when the AI reaches its limits.

This "opacity" of implementation necessitates a prioritization of "explainability" and "AI-based verification" (AI V&V). AI V&V involves the use of independent validation agents to evaluate the outputs of generation agents, ensuring the human orchestrator can trust the code's integrity [7]. Without these mechanisms, the risk of "systematic error" where an AI replicates the same logical flaw across a codebase remains high, eroding the foundational trust required for effective human-AI teaming.

## 5 Synthesizing Synergy Through Human-AI Orchestration

The evolution of the developer's role into that of an "Orchestrator" is essential for successful synergy. In this paradigm, humans manage, validate, and direct autonomous agents rather than focusing on raw syntax generation [7]. Several methodologies have emerged to facilitate this transition:

- **Agentsway:** Proposed by J. Eranga *et al.* (2025), this framework organizes agents into team structures with defined communication protocols, assigning roles such as "Code Manufacturer" to ensure reasoning remains dedicated to specific domains.
- **Agile Integration:** M. Usman (2025) developed a framework for integrating autonomous MAS into "Agile" cycles, allowing AI to assume roles such as "Scrum Master" or "Tester" within established sprint schedules.
- **Usability Patterns:** A. Soni (2025) identified context awareness and progressive complexity disclosure as critical patterns for maintaining developer agency without overwhelming cognitive capacity.
- **Intelligent Pair Programming:** R. Appachikumar (2025) redefines collaboration by replacing the human partner with an AI capable of providing real-time architectural guidance and bug identification.

Table 1: Comparative Analysis of 2025-2026 AI Agent Methodologies.

Methodology	Primary Focus	Human Role	Key Publication
Agentic SE Process	Automation	Quality Control	Otoum <i>et al.</i> (2026)
Agentsway Team Organisation	Team Organisation	Protocol Oversight	Eranga <i>et al.</i> (2025)
Agile Integration	Project Management	Sprint Management	Usman (2025)
HULA (HITL)	Feedback Loops	Decision Control	Roychoudhury <i>et al.</i> (2025)
Proactive Design	Interruption Management	Focus Maintenance	Pu <i>et al.</i> (2025)
Architectural Usability	UI Patterns	Context Control	Soni (2025)

The Hybrid Intelligence Model of Human-AI Teaming

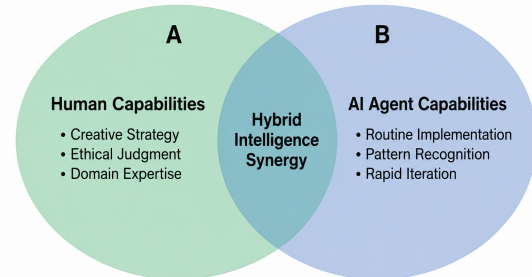


Figure 3: Hybrid Intelligence Venn Diagram showing overlap of human creativity and agentic automation.

## 6 Strategic Evaluation

Based on real-world viability and the 2025–2026 research landscape, methodologies are ranked as follows:

1. **Trust & Verification** (A. Roychoudhury *et al.*, 2025): The highest priority. Technical reliability and AI V&V are prerequisite for enterprise adoption; speed is irrelevant without safety.
2. **Agentsway / Organisation** (J. Eranga *et al.*, 2025): Essential for scaling agentic reasoning to complex, multi-tasking team projects.
3. **Proactive Design** (J. Pu *et al.*, 2025): Critical for balancing productivity with human focus through interruption management.

A significant gap persists between research prototypes and enterprise reality. Most academic tools, such as the initial versions of Codellaborator, focused on single-file Python environments, whereas real-world engineering involves multi-file interdependencies across massive codebases. Furthermore, the J. Pu *et al.* (2025) study utilized a sample size of only 18 participants, which, while deeply insightful regarding cognitive load, requires broader validation across large-scale industrial JIRA deployments. To bridge this, "Social Transparency" the explicit tracking of who or what modified code and for what reason must become a first-class design criterion to resolve the inherent opacity of LLM-based reasoning.

## 7 Conclusion

The trajectory of software engineering clearly indicates a transition toward Hybrid Intelligence. While proactive agents like the HULA framework can significantly bolster productivity, as evidenced by its 82% plan approval rate, they introduce systemic risks regarding workflow disruption and the erosion of technical depth. Success in this landscape requires developers to move beyond syntax generation to become strategic orchestrators of autonomous systems. Provided that rigorous standards of trust, transparency, and AI-based verification are maintained, agentic automation will expand human capabilities rather than replace them, fostering a more creative and efficient engineering profession.

## References

- [1] R. Appachikumar. Intelligent pair programming: Redefining collaboration between developers and ai agents. *European Economics Letters*, 13(3), 2025.
- [2] L. Chen et al. Code with me or for me? how increasing ai automation transforms developer workflows. *arXiv preprint arXiv:2507.08149*, 2025.
- [3] J. Eranga et al. Agentsway – software development methodology for ai agents-based teams. *arXiv preprint arXiv:2510.23664*, 2025.
- [4] S. Otoum, M. Al-Ayyoub, Y. Jararweh, and E. Benkhelifa. Methods and techniques of agentic software engineering: A systematic literature review. *IEEE Access*, 2026.
- [5] J. Pu et al. Assistance or disruption? exploring and evaluating the design and trade-offs of proactive ai programming support. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '25)*, 2025.
- [6] A. Roychoudhury. Agentic ai for software: Thoughts from software engineering community. *arXiv preprint arXiv:2508.1743*, 2025.
- [7] A. Roychoudhury, W. Takerngsaksiri, et al. Ai software engineer: Programming with trust. *arXiv preprint arXiv:2502.13767*, 2025.
- [8] A. Soni. Interactive developer tools powered by ai agents: Usability and architectural patterns. *European Modern Studies Journal*, 9(5), 2025.
- [9] M. Usman. Autonomous multi-agent llms in agile development: A framework for ai-driven collaboration. *Research-Gate*, 2025.
- [10] X. Wang, S. Gulwani, J. X. He, S. Jain, J. Kincaid, M. Maron, and O. Polozov. Ai agentic programming: A survey of techniques, challenges, and opportunities. *arXiv preprint arXiv:2508.11126*, 2025.