The Augmented Enterprise: A Per-Employee AI Agent Model for Organizational Transformation

Version: 1.0
Date: March 11, 2026
Author: Raúl Santamaría-Rivera
Contact: rs@raulsantamaria.com
Status: Preprint — Not Peer-Reviewed

Abstract

This paper introduces the concept of the **Augmented Enterprise**, an organizational paradigm in which every employee is paired with a dedicated, role-specific AI agent. Unlike existing models that focus on enterprise-wide AI solutions or general-purpose copilots, the Augmented Enterprise proposes a granular, human-centric approach — deploying AI intelligence at the individual level.

We establish a theoretical framework grounded in the Centaur model of human-AI collaboration, distinguishing between Augmented and Artificial Intelligence, and explore how this model reshapes organizational architecture from traditional hierarchies to dynamic networks of human-agent teams. The paper outlines an agent typology calibrated to distinct role categories and examines how early adopters are already demonstrating the viability of this approach in practice.

Our central argument is that the transformative unit of AI adoption is not the organization, not the team, but **the individual employee** — and that delivering AI value at that level requires an approach fundamentally different from current enterprise AI products.

What is needed is not an organizational AI layer but a personal AI agent for every employee: one that knows their role, their tools, their communication style, and their recurring responsibilities.

Table of Contents

---

1. Introduction

Knowledge work is undergoing its most significant structural transformation since the introduction of personal computing. For decades, productivity tools — email clients, spreadsheets, project management platforms — augmented individual capability by improving access to information and reducing manual effort. What they could not do was *act*: initiate follow-ups, synthesize context across systems, manage relationships, or coordinate multi-step tasks autonomously.

The emergence of capable large language models (LLMs) and agentic AI systems changes this calculus fundamentally. For the first time, it is technically and economically viable to deploy software that does not merely assist a human in performing a task, but *acts alongside them* — observing context, taking initiative, and executing on their behalf.

Yet the dominant paradigm for enterprise AI adoption in 2025-2026 remains horizontal: a single AI system deployed across an organization, accessible to all employees through a unified interface. Microsoft Copilot, Google Gemini for Workspace, and similar offerings treat AI as an organizational utility — available on demand, but generic by design.

This paper argues that such an approach fails to capture the majority of available value. The effective application of AI at the individual level requires configuration, context, and continuity — properties that cannot be delivered by a shared, stateless system. What is needed is not an organizational AI layer but a **personal AI agent for every employee**: one that knows their role, their tools, their communication style, and their recurring responsibilities.

We call this model the **Augmented Enterprise**. While the term "enterprise" evokes large corporations, the model is scale-agnostic by design. A five-person startup, a municipal government office, a nonprofit with twelve staff — any organization where knowledge workers could benefit from a dedicated AI partner is a candidate. The principles and architecture described in this paper apply wherever there are people whose work involves communication, coordination, and decision-making. The word "enterprise" here denotes not a size of organization but a mode of operation: deliberate, structured, and serious about leveraging AI at the individual level.

---

2. Theoretical Framework

2.1 The Centaur Model: Human-Algorithm Symbiosis

The concept of the Centaur — a hybrid human-algorithm entity that outperforms both humans and machines operating independently — has deep roots in decision science. Saghafian and Idan (2024) formalize this concept in the context of generative AI, defining centaurs as "hybrid human-algorithm models that combine both formal analytics and human intuition in a symbiotic manner within their learning and reasoning process."

The Centaur model is relevant to the Augmented Enterprise for a foundational reason: it establishes that the optimal unit of AI-assisted work is *the human-agent pair*, not the isolated human or the isolated machine. As Kasparov famously demonstrated in freestyle chess, "weak human plus machine plus better process was superior to a strong computer alone and, more remarkably, superior to a strong human plus machine plus inferior process" (Kasparov, 2010, as cited in Saghafian & Idan, 2024).

This insight scales to the enterprise context. An employee with a well-configured, role-aware AI

agent — one that knows their workflow, their contacts, their ongoing commitments — does not merely work faster. They work *qualitatively differently*, able to maintain more relationships, track more commitments, and act on more opportunities simultaneously than any unaided individual could.

The research on centaur models in clinical settings (Orfanoudaki et al., 2022) further confirms that this advantage is not marginal: centaur systems consistently outperform both solo experts and solo algorithms across complex decision-making domains. The implication for knowledge work is significant. Empirical evidence suggests that this model resonates beyond professional circles: in a study of 1,841 respondents aged 14–39, 27.3% already identified as "digital centaurs" — individuals who consciously integrate AI to augment their personal capabilities — with an additional 41.3% aspiring to adopt this strategy within the next decade (Soldatova et al., 2025).

Closely related is Randazzo et al.'s (2024) taxonomy in "Cyborgs, Centaurs and Self-Automators," which identifies three distinct modes of human-GenAI knowledge work:

- **Cyborgs:** AI deeply integrated into the human's cognitive process
- **Centaurs:** Human and AI operating in parallel, each handling distinct task components
- **Self-Automators:** Humans who use AI to fully automate tasks they previously performed manually

The Augmented Enterprise model supports all three modes. A dedicated agent can serve as a cognitive extension woven into the employee's thinking process (Cyborg), as a parallel partner handling distinct task components (Centaur), or as a delegated executor of fully automatable tasks (Self-Automator). The same agent may operate in different modes throughout a single workday depending on the task at hand. In Randazzo et al.'s framing, the question is not whether AI replaces human work, but *how* the human chooses to integrate AI into their cognitive process — and the Augmented Enterprise provides the infrastructure for that choice to be made deliberately, at scale.

2.2 Augmentation vs. Automation: A Critical Distinction

A persistent ambiguity in enterprise AI discourse confuses augmentation with automation. The distinction has both practical and strategic implications.

Automation replaces human activity: the machine performs a task that a human previously performed, eliminating the human from the loop. This is the logic of robotic process automation (RPA), most customer service bots, and document processing pipelines.

Augmentation expands human capacity: the human remains the primary actor, but their effective reach, speed, and cognitive bandwidth is multiplied by AI assistance. As conceptualized by IEEE and Gartner to distinguish collaborative AI from fully autonomous systems, and formalized by Fügener, Walzner & Gupta (2025) in *Management Science*, augmented intelligence preserves human agency while extending its scope. Chiriatti et al. (2024) sharpen this idea with the concept of "System 0" — a cognitive layer in which AI acts as a preprocessor that filters, organizes, and structures information *before* the human engages either intuitive (System 1) or analytical (System 2) thinking. In the Augmented Enterprise, the agent operates across all three layers. As System 0, it monitors channels, classifies inputs, tracks commitments, and presents an already-digested view of the employee's operational landscape. But the agent also operates at the System 2 level: it synthesizes information from multiple sources, drafts analytical reports, compares options, and researches questions that would otherwise consume hours of deliberative human effort. The result is

not merely faster access to information but an expansion of the employee's analytical capacity itself — the human thinks better because the agent has already done much of the heavy cognitive lifting.

The distinction matters for the Augmented Enterprise because its value proposition is *capability multiplication*. The goal is not to replace the salesperson with an AI but to give every salesperson the follow-up discipline, analytical depth, and responsiveness of a top performer. Not to eliminate the project manager but to ensure they never drop a thread, miss a deadline signal, or fail to synthesize a status update.

This has direct implications for adoption. Employees — the actual users — are more willing to embrace tools that enhance their performance than those that threaten their roles. An AI agent positioned as a personal operational partner encounters less resistance than one positioned as a replacement for human labor, regardless of how the organization frames it. The Augmented Enterprise is designed around this reality: agents succeed because employees want to use them, not because management mandates it.

2.3 Emerging Organizational Paradigms

The Augmented Enterprise model did not emerge in a vacuum. Several converging frameworks from leading research and consulting institutions point toward similar conclusions.

The Frontier Firm (Microsoft, 2025)

Microsoft's Work Trend Index Annual Report (2025) describes the emergence of "Frontier Firms" — organizations that blend machine intelligence with human judgment through hybrid teams of humans and agents. Microsoft identifies three phases of evolution: AI as assistant, AI as digital colleague, and AI as autonomous process executor. The Frontier Firm operates across all three phases simultaneously, deploying agents that are "AI-operated but human-led."

The Microsoft data is striking in its urgency: 82% of leaders consider 2025 a pivotal year for rethinking strategy and operations, and 81% expect agents to be extensively integrated within 12-18 months. More tellingly, 53% of leaders report that productivity must increase while 80% of the workforce reports lacking sufficient time and energy to perform their work. The Frontier Firm bridges this gap through digital labor.

The Agentic Organization (McKinsey, 2025)

McKinsey's framework (2025) identifies the "Agentic Organization" as the next paradigm for the AI era, characterized by systems that "plan, act, and learn on their own." Their research highlights that 76% of executive survey respondents view agentic AI as more like a coworker than a tool — a perception shift with profound implications for deployment strategy.

McKinsey identifies four operational tensions in adopting agentic AI: the tool-coworker duality, the need to simultaneously manage AI as both capital asset and workforce, competing demands from IT, finance, HR, and business leadership, and the challenge of redesigning processes before strategy is formalized. Notably, 35% of organizations have already deployed agentic AI despite having no strategic framework in place — a gap the Augmented Enterprise model is positioned to address.

The Agentic Enterprise (MIT Sloan, 2025)

MIT Sloan's research on the Emerging Agentic Enterprise frames the core challenge as a management problem: agentic AI "simultaneously influences the design of processes, the structure of roles, the allocation of decision rights, and the culture of accountability." The separation between technology

executives and strategic executives — historically functional — becomes untenable when AI is both a tool and a colleague.

The MIT Sloan framework is particularly relevant to the Augmented Enterprise model because it emphasizes the organizational redesign required to realize value from agentic AI. Deploying per-employee agents is not merely a technology decision — it requires rethinking how roles are defined, how autonomy is delegated, and how human-agent accountability is structured.

SHINE and the Human Operating System (Chief Learning Officer, 2025)

Salesforce's concept of the "Agentic Enterprise" and the associated SHINE framework (Skill, Heart, Intuition, Networks, Ethics) emphasizes that agentic AI transformation succeeds only when the human operating system is as capable as the AI layer. This reinforces the Augmented Enterprise's human-centric design philosophy: agents are configured to serve specific humans, not to replace them.

---

3. The Augmented Enterprise: Core Concept

3.1 Definition and Principles

The Augmented Enterprise is an organization in which every employee is paired with a dedicated, role-specific AI agent that acts as their personal operational partner.

This definition carries four critical qualifiers:

1. **Every employee** — not selected power users, not a shared organizational tool. The value of the model grows exponentially with coverage: an organization where all employees are augmented achieves network effects that partial deployment cannot.
2. **Dedicated** — each agent is configured for a specific individual, not a generic AI accessible to all. The agent knows the employee's role, their context, their communication patterns, and their recurring responsibilities.
3. **Role-specific** — configuration is determined by what the employee actually does, not by their seniority. A junior salesperson and a senior analyst have different agent configurations because their work is different.
4. **Operational partner** — the agent acts, not merely advises. It manages follow-ups, drafts communications, researches questions, coordinates with other systems, and executes tasks autonomously within defined boundaries.

Core principles:

- **Presence in natural channels:** Agents operate where employees already work — Microsoft Teams, Slack, WhatsApp, Telegram, email — not in a separate AI interface requiring behavioral change.
- **Conversational interface:** The interaction between employee and agent is natural language conversation — the same medium humans already use to coordinate with colleagues. This eliminates the learning curve inherent in traditional software interfaces. The employee does not learn a new tool; they talk to a partner — by text or by voice, with speech recognition and text-to-speech enabling hands-free interaction when needed. This extends the model's reach beyond desk-bound knowledge workers to employees in the field, in transit, or in environments

where typing is impractical. Each conversation enriches the agent's understanding of the employee's context, creating a compounding effect: the more you work with your agent, the better it works with you.

- **Continuity of context:** The agent maintains memory across interactions, building a persistent model of the employee's work, relationships, and commitments.
- **Managed evolution:** Agents are maintained, improved, and adapted over time — treating agent capability as an evolving asset rather than a static deployment.
- **Technical abstraction:** The employee interacts with a conversational partner, not with a software system. All underlying complexity — model selection, tool integration, scheduling, error recovery, credential management — is invisible to the user. The employee does not configure, maintain, or troubleshoot their agent; a supervisory layer (described in 3.3) handles this autonomously. This is a deliberate design choice: if using the agent requires technical literacy, the model fails to reach the majority of knowledge workers it is designed to serve.
- **Human primacy:** All consequential actions are flagged for human approval. The agent amplifies human judgment; it does not replace it.

3.2 The Per-Employee Agent Model

The unit of value in the Augmented Enterprise is the human-agent pair, not the individual employee or the isolated agent. This distinction is not semantic — it has architectural implications.

Each human-agent pair operates as a micro-unit with its own:

- **Context window:** persistent memory of the employee's work, commitments, and relationships
- **Tool surface:** access to the specific systems and channels the employee uses
- **Communication persona:** calibrated to the employee's style, tone, and organizational role
- **Autonomy envelope:** defined scope of independent action, with escalation protocols for edge cases

The agent does not operate in isolation. Within the Augmented Enterprise, agents form a network — capable of coordination, information sharing, and task handoffs — while each remains accountable to a specific human. This mirrors the human organizational structure: individuals who collaborate through defined channels and hierarchies, each with their own role and responsibilities.

Agent configuration varies with role complexity. Not all agents are identical. Configuration differs along three primary dimensions: autonomy (how much the agent can do without asking), tool access (which systems the agent can reach), and coordination capability (whether the agent can delegate to other agents).

At one end of the spectrum, an agent paired with an administrative assistant operates within tight boundaries: it drafts emails, tracks follow-ups, manages calendars, and surfaces relevant information — but escalates anything ambiguous to its human. At the other end, an agent paired with a department director may orchestrate multi-step research by spawning temporary sub-agents, synthesize inputs from across the organization, and initiate actions across multiple systems with minimal supervision.

Between these extremes lies a continuum calibrated to what each person actually does. A project manager's agent maintains relationship context across dozens of stakeholders, tracks commitments and deadlines, and prepares briefings by pulling from multiple sources. A technical lead's agent reviews code, analyzes architecture decisions, coordinates development tasks across repositories, and leverages specialized models for different aspects of software engineering.

The key insight is that the same underlying architecture supports all of these configurations. What changes is not the system but the parameters: the scope of autonomous action, the tools connected, the escalation thresholds, and the communication persona. This makes the model both scalable and adaptable — a new employee can be paired with an appropriately configured agent without redesigning the platform.

3.3 The Supervisory Layer

The per-employee agent model described in 3.2 creates a distributed network of human-agent pairs, each configured for a specific individual and role. But a distributed system of autonomous agents introduces an operational challenge that traditional software does not face: who maintains the agents themselves?

In conventional enterprise software, the answer is straightforward — an IT department handles updates, monitors system health, and resolves technical issues. But the Augmented Enterprise is not conventional software. Each agent operates with its own persistent memory, its own tool integrations, its own communication channels, and its own evolving configuration. When something breaks — a credential expires, a scheduled task fails, a communication protocol between agents falls out of sync — the affected employee should not need to diagnose or repair it. The entire premise of the model is that the agent removes operational complexity from the employee's work, not that it introduces a new category of technical burden.

This is the role of the **Supervisory Layer**: an architectural tier that sits above the network of per-employee agents and is responsible for the health, continuity, and self-repair of the system as a whole.

The Operational Oversight Agent. At the center of the Supervisory Layer is a dedicated agent — the Operational Oversight Agent (OOA) — whose "employee" is the system itself. Unlike per-employee agents, the OOA is not paired with a knowledge worker performing business tasks. Its role is infrastructure stewardship: it monitors the operational health of all other agents, detects failures before they become visible to users, and resolves technical issues autonomously or with minimal human intervention.

The OOA's responsibilities include:

- **Proactive system monitoring.** Continuous review of agent health, scheduled task execution, communication channels, credential validity, and resource consumption. Issues are detected and addressed before they affect the employee experience.
- **Self-repair and autonomous remediation.** When an agent encounters a technical failure — a broken integration, a stale authentication token, a misconfigured schedule — the OOA diagnoses the root cause and applies the fix. The affected agent's user may never be aware that anything went wrong.
- **Inter-agent technical support.** Per-employee agents are configured for business tasks, not for debugging infrastructure. When an agent encounters a problem outside its operational scope, it escalates to the OOA through internal protocols — transparently to the user. The OOA resolves the issue and returns control to the requesting agent, which continues its work seamlessly.
- **Configuration management.** As the system evolves — new agents are onboarded, existing configurations are refined, communication protocols are updated — the OOA ensures consistency across the network. It maintains the system's documented architecture, validates

changes, and prevents configuration drift.

- **Escalation to human administrators.** Not every problem can be resolved autonomously. The OOA maintains defined escalation thresholds: issues that exceed its autonomy envelope are reported to a human administrator with a clear diagnosis, the actions already attempted, and a recommended course of action. The human decides; the OOA executes.

Why this matters for adoption. The Supervisory Layer addresses one of the most significant barriers to enterprise AI deployment: the assumption that AI systems require technical staff to maintain. In the Augmented Enterprise, the system maintains itself. An employee whose agent encounters a credential expiration does not file a support ticket — the OOA detects the issue, renews the credential or alerts an administrator, and the agent resumes normal operation. A manager whose agent's scheduled briefing fails to generate does not troubleshoot a cron job — the OOA identifies the failure, repairs it, and ensures the next briefing arrives on time.

This self-maintaining capability transforms the operational economics of the model. Without a Supervisory Layer, the cost of maintaining N agents scales linearly with N — each additional agent requires proportional technical oversight. With an OOA, the marginal cost of adding a new agent approaches the cost of configuration alone, since ongoing maintenance is handled by the system itself.

The OOA as an emergent property of the model. In practice, the Supervisory Layer often emerges naturally rather than being designed top-down. In the author's operational deployment, the personal operational agent — originally configured for business tasks like email management and scheduling — gradually assumed infrastructure responsibilities as the system grew. It began monitoring other agents' health, repairing failed tasks, and coordinating technical changes across the network. This evolution was driven by necessity: as the number of agents increased, manual maintenance became unsustainable, and the most capable agent in the system absorbed the role. The Augmented Enterprise, in this sense, does not just augment employees — it augments itself.

3.4 Organizational Transformation

The deployment of per-employee agents does not only accelerate existing workflows. It transforms the organizational architecture in several aspects.

Effective capacity multiplication. A well-configured agent pair effectively doubles an employee's operational bandwidth for the tasks the agent handles. An organization of ten augmented employees operates with the effective capacity of fifteen to twenty, depending on role composition and agent capability.

Reduction of cognitive load. A significant portion of knowledge work involves tracking — who owes what to whom, what follows what, what was decided and when. Agents absorb this tracking function entirely, freeing human attention for higher-order work.

Democratization of elite capabilities. Top performers in any role share common traits: rigorous follow-up, proactive communication, contextual analysis before action. Agents codify these behaviors and make them available to every employee, regardless of experience level. Experimental evidence supports this effect: in a study of 758 consultants, AI-augmented performance improved by 43% among below-average performers compared to 17% for above-average ones (Dell'Acqua et al., 2023, as cited in Soldatova et al., 2025). The augmented enterprise does not raise the ceiling — it raises the floor.

Acceleration of organizational learning. As agents accumulate context about an organization's workflows, relationships, and decision patterns, they become repositories of institutional knowledge — more durable than any individual employee's memory and more accessible than any documentation system.

Early evidence. Early adopters are already demonstrating capability multiplication at its most extreme. Solo practitioners using orchestrated agent systems have reported operating at the output level of small teams — managing dozens of daily deliverables without directly performing the underlying tasks (Sun, 2026). While these are individual cases rather than organizational deployments, they validate the core mechanism: a human with well-configured agents produces qualitatively more than a human with conventional tools.

The author's own deployment provides a more direct illustration. A multi-agent system has been operational since early 2026, comprising a personal operational agent (managing email, follow-ups, scheduling, and system administration), a technical lead agent (coordinating software development through subordinate coding agents), and dedicated agents assigned to different external clients — coordinating through structured protocols and communicating through standard messaging channels. This system operates continuously, with agents maintaining persistent memory and escalating to humans only when their autonomy envelope is exceeded. It serves simultaneously as a research platform and as the operational backbone of a consultancy practice — the Augmented Enterprise applied to itself.

---

4. Open Challenges and Future Work

Several aspects of the Augmented Enterprise model require further development before large-scale deployment.

Legal and regulatory compliance. Processing employee data on behalf of organizations creates obligations under GDPR and similar frameworks. Data residency, retention policies, and breach notification procedures must be defined. The boundaries of agent autonomy — what an agent may do without explicit human approval — have implications that existing regulatory frameworks do not fully address.

Integration complexity. The value of an agent is directly proportional to the depth of its tool integration. An agent with access only to a messaging channel delivers a fraction of the value of one integrated with email, calendar, CRM, and project management. Each integration represents engineering effort and introduces potential failure modes.

Quality assurance and error handling. Agents acting on behalf of employees will make errors. Some errors will be low-stakes (a poorly phrased draft); others will be consequential (a follow-up sent to the wrong contact, a commitment made without authorization). Robust error detection, escalation, and remediation protocols are prerequisites for organizational deployment.

Governance and accountability. When an agent acts on an employee's behalf, who is accountable for the outcome? As MIT Sloan Management Review (2025) identifies, this is a central challenge of the agentic enterprise. The Augmented Enterprise must define clear accountability models that preserve human responsibility while enabling agent autonomy.

Organizational change management. Deploying per-employee agents is not a technology rollout — it is an organizational transformation. Employees must understand what their agent can and

cannot do, trust its capabilities, and adapt their workflows to leverage the partnership. This requires structured onboarding, ongoing training, and cultural adaptation. Beyond the operational dimension, the behavioral impact on employees deserves attention: how the introduction of an AI partner affects role perception, sense of autonomy, team dynamics, and power structures within the organization. Early research on agentic AI in organizational behavior suggests that these psychological and relational effects may be as important as the productivity gains (Joshi, 2026).

Measurement and evaluation. How do you measure the ROI of augmentation? Traditional productivity metrics (output per hour, tasks completed) capture only part of the value. The deeper effects — reduced cognitive load, improved follow-through, better relationship maintenance — are harder to quantify but may represent the majority of the benefit.

Scalability of configuration. The current model assumes high-touch configuration of each agent for each employee. This is appropriate for early deployments but does not scale without tooling. Role-based templates, automated configuration pipelines, and self-adapting agents are necessary for deployment beyond small organizations.

---

5. Conclusion

The Augmented Enterprise represents a specific and actionable answer to a broadly recognized challenge: how to translate the general capability of AI into specific, measurable value for individual knowledge workers.

Existing enterprise AI products are built around the logic of organizational access — make AI available to the whole company and let each employee find their own use. This approach produces uneven adoption, limited depth of use, and no lasting advantage for the organizations that deploy it.

The Augmented Enterprise inverts this logic. Rather than making AI available to employees, it assigns AI to employees — role-calibrated, context-aware agents that operate as persistent partners in the natural channels where work already happens.

The theoretical basis for this model is well established: the Centaur framework's evidence that human-agent pairs outperform both humans and machines operating alone; the Frontier Firm's data on the capacity gap that digital labor is positioned to close; the Agentic Organization's insight that agentic AI demands new organizational logic, not just new tools. Early practitioners and operational pilots are already demonstrating the viability of the model, with results that suggest significant organizational impact when deployed at scale.

What is missing from the current landscape is not the technology — it is the organizational framework for deploying it at the granularity where it creates the most value: the individual employee. This paper proposes that framework.

The Augmented Enterprise is not a future possibility. It is operational today. What remains is the willingness to rethink not just what tools employees use, but what kind of partner they work with.

References

Randazzo, S., Lifshitz-Assaf, H., Kellogg, K. C., Dell'Acqua, F., Mollick, E. R., Candelon, F., & Lakhani, K. R. (2024). *Cyborgs, Centaurs and Self-Automators: The Three Modes of Human-GenAI*

Knowledge Work and Their Implications for Skilling and the Future of Expertise. Harvard Business School Working Paper No. 25-022. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4921696

Kasparov, G. (2010). The chess master and the computer. *The New York Review of Books, 57*(2), 16–19.

McKinsey & Company. (2025). *The agentic organization: Contours of the next paradigm for the AI era*. McKinsey Global Institute. https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-agentic-organization-contours-of-the-next-paradigm-for-the-ai-era

Microsoft. (2025). *2025 Work Trend Index Annual Report: The year the Frontier Firm is born*. Microsoft WorkLab. https://www.microsoft.com/en-us/worklab/work-trend-index/2025-the-year-the-frontier-firm-is-born

MIT Sloan Management Review. (2025). *The Emerging Agentic Enterprise: How Leaders Must Navigate a New Age of AI*. https://sloanreview.mit.edu/projects/the-emerging-agentic-enterprise-how-leaders-must-navigate-a-new-age-of-ai/

Orfanoudaki, A., Saghafian, S., Song, K., Chakkera, H. A., & Cook, C. B. (2022). *Algorithm, human, or the centaur: How to enhance clinical care?* HKS Faculty Research Working Paper Series RWP22-027. Harvard Kennedy School.

Fügener, A., Walzner, D. D., & Gupta, A. (2025). Roles of artificial intelligence in collaboration with humans: Automation, augmentation, and the future of work. *Management Science*. https://pubsonline.informs.org/doi/10.1287/mnsc.2024.05684

Saghafian, S., & Idan, L. (2024). Effective generative AI: The human-algorithm centaur. *Harvard Data Science Review*, Special Issue 5. https://hdsr.mitpress.mit.edu/pub/3rvlzjtw

Salesforce / Chief Learning Officer. (2025). *SHINE in the age of agentic AI: The human operating system behind enterprise transformation*. https://www.chieflearningofficer.com/2025/12/29/shine-in-the-age-of-agentic-ai-the-human-operating-system-behind-enterprise-transformation/

Sun, E. (2026). *OpenClaw + Codex/Claude Code Agent Swarm: The One-Person Dev Team*. X (formerly Twitter). https://x.com/elvissun/status/2025920521871716562

Joshi, S. (2026). Agentic generative artificial intelligence in enterprise organizational behavior: An integrated scholarly-practitioner mathematical and theoretical framework. *Preprints*. https://www.preprints.org/manuscript/202602.0148

Chiriatti, M., Ganapini, M., Panai, E., Ubiali, M., & Riva, G. (2024). The case for human–AI interaction as system 0 thinking. *Nature Human Behaviour*, 8(10), 1829–1830. https://doi.org/10.1038/s41562-024-01995-5

Soldatova, G. U., Rasskazova, E. I., & Chigarkova, S. V. (2025). The digital centaur as a type of technologically augmented human in the AI era: Personal and digital predictors. *Behavioral Sciences*, 15(11), 1487. https://pmc.ncbi.nlm.nih.gov/articles/PMC12649359/

Preprint — this paper has not undergone peer review. Comments welcome at rs@raulsantamaria.com.