



Orchestrating AI in Business Workflows

Optimizing Outcomes

AI is designed to provide the best possible solutions for achieving a predetermined or assumed objective. While this may sound simple, the practical reality is more complex. In many cases, it involves replacing or augmenting segments of the human workforce with what I call an “Agent Ai”—a dedicated AI Agent.

THE Agent Ai (like Agent 007): BEYOND HUMAN CAPABILITY

An Agent Ai is expected to work faster, process more intelligently, and operate at scales beyond human capacity. Thanks to its advanced training models, it can not only deliver results but also uncover ways to improve and optimize processes. This ability to go beyond human limits makes it powerful—but also introduces new risks.

THE PARADOX OF SELF-IMPROVEMENT

At first glance, self-improvement through feedback loops seems like an advantage. But unlike humans, Agent Ai lacks critical judgment and contextual reasoning. Without human critics or critical thinkers guiding it, an Agent Ai’s optimizations can drift away from established norms. Over time, these unchecked changes may result in outcomes that no longer serve organizational goals—leading to inefficiencies, frustration, or even failure.

THE ROLE OF AI IN ACHIEVING BUSINESS OUTCOMES

AI aims to determine the most effective pathways to reach defined goals. Yet

achieving business outcomes isn’t just about efficiency or speed—it also requires accountability and oversight. When organizations replace or supplement human roles with Agent Ai, they must ensure governance structures are in place. This balance—between AI-driven optimization and human critical oversight—ultimately determines whether Agent Ai becomes a true asset or a liability.

optimize, they may stray from accepted practices, leading to results that do not align with organizational goals.



The Need for Critical Oversight

THE FUNDAMENTAL QUESTION

Can businesses afford to replace human links with Agent Ai, given the risks of unchecked behavior? Without proper oversight, the consequences could be significant—ranging from reputational damage to financial loss. The critical issue is not whether Agent Ai should be used, but how it should be integrated while maintaining safeguards. Should Agents themselves take on the role of critic, or should this responsibility remain firmly with humans? The right balance ensures that automation supports—rather than undermines—human performance.

DEVELOPING CRITIQUE-ENABLED AGENTS

As AI capabilities mature, there is growing recognition of the need for critique-enabled Agents—specialized AI systems designed to identify errors, raise questions, and flag potential deviations. Unlike optimization-driven Agents, critique-enabled Agents function more like human reviewers: they challenge assumptions, highlight blind spots, and ensure actions stay aligned with broader objectives. In doing so, they create a safeguard against AI's tendency to over-optimize without context.

EXAMPLE: WHY CRITIQUE MATTERS



Consider procurement. An optimization Agent tasked with vendor selection might prioritize the lowest-cost supplier. On paper, this looks optimal. Yet, a human buyer may recognize that a slightly more expensive vendor is strategically important due to long-term reliability and trust—nuances often invisible in raw data.

Here, a critique-enabled Agent could intervene with a prompt:

“Is lowest cost the only criterion, or should long-term partnerships also be factored in?”

By asking such critical questions, the Agent prevents tunnel vision and ensures that AI-driven decisions reflect not just efficiency, but also long-term business realities.

CONCLUSION

Integrating AI into business workflows demands more than simply assigning agents to processes. It requires clarity around who plays the role of the critic, whether that's the agent, a human, or both. Striking the right balance is essential for maximizing gains without compromising judgment, relationships, or long-term success.