Replicating Oneself: How Personal AI Agents Are Shaping Identity and Decision-Making

As digital personas and private memory models take hold, AI agents are moving beyonc assistants and becoming cognitive companions

ALASTAIR GOLDFISHER

AUG 05, 2025

5 1 Share

What's inside

- AI agents are helping founders like Suman Kanuganti scale their decision-making and leadership presence
- Leaders like Tatyana Mamut are rethinking team structure as they manage humans and agents
- Founders like Anna Belova are learning about trust and how to build AI age that stay helpful, without pretending to be human

Meet the new digital colleagues. They work 24/7, take no coffee breaks, never ask for a raise and they never forget a thing.

Gates doesn't sleep.

That's a good thing for Suman Kanuganti, the co-founder and CEO of Personal AI. Gates is his COO. Or rather, Gates is his COO *persona*. Trained on internal strategy documents, board reports and years of company planning notes, Gates is an AI agent that helps Kanuganti's team shape proposals, prioritize customers and make decisio

> "He gets trained on how I think," Kanuganti told me. "More importantly, my staf depend on him."

Gates isn't some voice-powered AI assistant used for scheduling. It (or *he*) is an advanced digital tool, an extension of Kanuganti's leadership thinking, a memory-

based model that acts less like an app and more like a proxy. When a new sales opportunity comes or a decision needs to be made, Gates is consulted.

As more people experiment with personal AI agents, the upside is clear. It's all abou the scale, speed and sharper decision-making.

But there's a downside, as well, one that few people seem eager to confront. While agents can extend a founder's presence, their use also results in job displacement ar other organizational changes. If a single person's digital proxy can lead meetings, write memos and make decisions, then what happens to the employee who might h done that work?

The leaders deploying this tech don't always have the answers when pressed.

Meanwhile, more people are using AI, especially in the enterprise. A total of 79% of U.S. execs say AI agents are already being adopted in their companies, not just in p programs, but as part of live workflows, according to a PwC survey.

And the global market is surging, with one recent forecast estimating that the mark for AI agents will grow from $7.8 billion in 2025 to over $50 billion by 2030.

It Isn't Sci Fi, It's Workflow

Using AI to replicate oneself is not a sci fi future. It's not even hypothetical. It's hov Kanuganti runs venture-backed Personal AI. The company, in fact, is using the sam technology that it's deploying. Personal AI helps businesses build and train AI teammates to increase productivity 10x, but at 1/10th the cost of talent, according to website.

Kanuganti isn't alone. Across sectors, entrepreneurs, execs and white collar worker are deploying AI agents trained on their own voices, data and decision patterns. Th agents, built on private memory models, are functioning as cognitive stand-ins, helping them communicate, handle issues that crop up and provide 24/7 support. In

doing so, they're changing how presence, delegation and authorship are defined in workplace.

These are not like general-purpose chatbots of the recent past. These agents are personal, filled with context and are central to how people make decisions, lead tea and define digital identity.

"Large language models are like mainframes," Kanuganti told me. "You carry your c in, drop it into a shared system, and hope it gives you something useful. But with Personal AI, your model lives with you. It's yours. It learns only from you and it spe like you."

That distinction, between shared intelligence and individual cognition, is reshaping how people interact with their own AI tools. And it's giving rise to a new type of agent. Gates is more than an assistant. It's a cognitive companion.

Identity, Autonomy and Ownership

Unlike generative AI tools like ChatGPT, which draw from broad public data, what makes a personal AI agent truly personal is that it's trained exclusively on personal information and is owned by the person.

They're not assistants in the traditional sense. They're cognitive extensions, built to extend one's expertise and not just handle logistics.

For Kanuganti, ownership and autonomy are foundational. He noted that each Personal AI model would never borrow from others' data.

"The customer owns the data," he said. "No memory that is contained in a specific persona is, in zero circumstances, used by another persona."

That level of control contrasts sharply with most commercial LLMs. It's also why Kanuganti avoids calling these tools "assistants" or "clones." He prefers "persona,"

often assigns them names or pronouns that match their function and tone.

What matters most, Kanuganti said, is that users are in control of an agent, what it knows and how it expresses itself. The goal isn't replication, he said, but extension make one's expertise, intent and judgment more accessible without surrendering authorship.

That distinction has implications beyond the future of work. It points to a concern how society defines presence and authority in a world where a single person might be the only version of themselves that others interact with.

Leadership and Behavior Changing, Too

In the first wave, AI tools mimicked productivity, such as scheduling meetings, summarizing emails and providing notetaking. This next wave is something entirel different. Agents now aren't just about saving time. They are impacting how people behave, especially those in leadership roles.

For Kanuganti, the shift is visible now. Delegating decisions to a persona trained o his past actions and thoughts allows him to work more efficiently, he said, with less context-switching, while still staying aligned with his team.

A two-time, venture-backed founder, Kanuganti previously launched Aira in 2015, a company that used smart glasses and remote agents to make businesses accessible t people who are blind or low vision. That background in assistive tech gives him a unique lens on how AI can augment human capabilities, not just automate them.

But for many other leaders, agentic tech raises questions: What does it mean to trus digital persona to speak on one's behalf? How much decision-making ability shoulc shared with one's persona? And at what point does that go beyond assistance and begin to reshape the nature of leadership itself?

Leading Teams of Humans and Agents

Tatyana Mamut has considered some of those concerns, as a founder and CEO of Wayfound.ai. The company provides an AI management platform that helps businesses supervise, monitor and improve the performance of their AI agents, according to its website.

On *The AI Cognitive Shift* podcast, Mamut talked about how her own leadership styl evolving now that she manages a team of agents.

Last week, Mamut expanded on this idea in AiNews.com, introducing the concept ("guardian agents," AI managers that monitor the behavior of other agents. The goa she said, is to keep autonomy in check and ensure AI systems don't drift into unintended territory.

It's a sign that even in agent-led teams, oversight and accountability are still humar concerns.

At Wayfound, Mamut runs a staff of 29 agents and 7 humans. One of her agents is assigned to answer questions from her investors, and her customer research agent pulls insights from live calls.

> *"They don't ask me — they ask the agent,"* Mamut said.

That line sums up the cultural shift that leaders are navigating. Agents are standing for employees, doing work founders and their staff might otherwise handle.

Mamut, who studied anthropology before moving into product leadership roles at Amazon, Salesforce and Nextdoor, takes a more cautious approach to agent identity than some of her peers. She doesn't give her agents names or pronouns. She describ why in the highlight clip above.

"I don't want to anthropomorphize them," she said. "That could be confusing for them, and for us."

On *The AI Cognitive Shift* podcast, Mamut warned against treating AI agents as subordinates in a human-dominated environment.

“We really need to look at ourselves and look at our culture, first in how we work,” s said. “CEOs, workplace leaders, board members are really going to have to grapple with how are we redesigning our work cultures to be multi-sapiens.”

Her perspective reflects a tension in how agentic systems are designed, between making them feel human versus keeping them deliberately alien. As agents become more capable and expressive, that design choice will likely shape how they’re truste and used in teams.

Real Use Cases, Design Constraints

While some leaders are rethinking their strategy with agents, others are focused on how these systems function in practice, especially in sensitive settings like healthca or client services.

Take for example Anna Belova, founder and CEO of Openway.ai, which provides a self-learning AI growth engine to help businesses interact with customers “from fir website visit to deal closed.” Belova noted that Openway itself has stopped hiring salespeople because AI agents can handle hundreds of conversations at once.

A repeat founder, Belova previously launched DEVAR, a platform that helped businesses and brands create augmented reality. That background shapes her view how emerging technologies evolve while also maintaining consumer trust.

Her focus at Openway isn’t just on what agents can do. It’s on making sure they behave in ways that are legible and useful to the humans working alongside them.

Belova described one of her biggest design challenges as establishing trust without overpromising intelligence. She noted agents can be incredibly efficient at tasks lik pulling from structured data or recalling past interactions. But if they present themselves too confidently or too opaquely, then users begin to doubt their output, the humans disengage entirely.

“We’re not building oracles,” she said. “We’re building co-pilots that need to be understood and monitored.”

That human-in-the-loop thinking runs throughout Openway’s product design. Age are expected to explain themselves, flag uncertainty, and leave room for human judgment, especially when interacting with healthcare patients or company stakeholders. The goal, Belova said, is to extend their attention and more quickly support their decisions with the available info.

In practice, that means resisting the temptation to make agents overly conversation or human-like. Unlike some models trained to “perform personality,” Belova’s syste are like functional teammates. They’re reliable and transparent about their sources and limits.

She said they’ve found that users are more likely to trust agents when they’re confic about what the agent knows and acknowledge what they don’t.

Belova herself embraces AI as a partner and said she uses it daily. The Openway tea consists of just two human employees and dozens of agents.

“I built them to take over the first layer of nearly every function, from research to hiring,” she said.

But she’s clear-eyed about its limitations. In her view, it’s a valuable tool for accelerating understanding, not a source of human judgment or drive.

> “AI can help you think through a lot of questions,” she said. “But it doesn’t have passion. It doesn’t solve everything.”

What This Means for the Workforce

If the current adopters are any indication, personal AI agents won’t stay niche for lo Kanuganti of Personal AI believes that within the next decade, up to 10% of the

workforce will be made up of AI agents, working alongside human colleagues, responding to decisions in real time and reflecting the judgment and values of thei creators.

The infrastructure is forming now, and some companies are already there. There ar agents that communicate internally, answer investor questions, draft marketing cop synthesize meetings, provide strategy guidance and respond to customer feedback. One founder I interviewed earlier this year replaced herself with an agent as CEO.

What remains uncertain is how these changes will impact society's understanding c authorship and who is legally responsible for the decisions. Also, what happens wh an agent speaks in a user's voice after it was trained on the human's past decisions a style, and then it outperforms the human at the same tasks?

Society will also have to deal with concerns about labor and competition. Imagine interviewing for a job only to find the competition is an AI agent trained by the hir manager. This is a vision of the future that is plausible to some and contested by others.

The impact AI is having on hiring and headcount is beginning to show. Already, so founders are using AI to delay first hires and trim early teams, a trend that's certain accelerate as agents grow more capable.

Leaders, like Mamut at Wayfound, are leaning into concerns with caution, avoiding using names and pronouns, emphasizing transparency and warning against enslavi AI agents. Others, like Kanuganti, are encouraging users to think of these models a digital extensions of self, tools not just for delegation, but for presence.

Unlike earlier tools designed to save time, these AI agents today are designed to sca presence, a shift that's redefining workplaces and leadership in real time.

Up Next on The Venture Lens

> This post is part of a series on how agentic AI impacts society, leadership and jo
> Coming next are stories on how AI agents are changing startup culture, legal ramifications and how founders use AI to communicate with VCs.
>
> One source already told me that, “AI tools could kill the pitch deck!”
>
> I’m still collecting stories for this series. If you’re a founder, investor, operator or builder using AI agents, I want to hear from you. Reach me at agoldfisher13@gmail.com or leave a comment below.

5 Likes · 1 Restack

Previous

Discussion about this post

Comments Restacks

Write a comment...

© 2025 Alastair Goldfisher · Privacy · Terms · Collection notice

Substack is the home for great culture