

AI Agent Guide: 800 Words From 8 Leaders

How today's innovators are redefining automation,
intelligence, and the future of work



Sanjay Srivastava
Chief Digital Strategist at
Genpact



Robb Wilson
CEO & Co-founder of
OneReach.ai



Leon Gordon
Current Microsoft Data Platform
MVP & Founder of Onyx Data



Douwe Kiela
CEO & co-founder of
Contextual AI



Surojit Chatterjee
Founder & CEO of Ema



Marianna Bonanome
Head of AI Strategy &
Partnerships at SandboxAQ



Chris Deotte
Senior Data Scientist at
NVIDIA



Andriy Burkov
Author of The Hundred-Page
Machine Learning Book



AI agents: Magic or mirage?

Every few years, a new “next big thing” in AI promises to change how we work. Most fall short.

Today, AI agents sit at that same inflection point: half headline, half hard reality.

Some executives are already automating customer service, research, and analytics with agentic systems. Others remain unconvinced that these models are truly intelligent or ready for business-critical tasks.

The truth, according to the eight leaders featured in this report, is more nuanced: AI agents are neither magic nor mirage.

Inside, you’ll discover:

- What is the scope of their role, and how could AI support it?
- How familiar are they with AI concepts and terms?
- What AI tools are they already using, and how confident are they?
- What skills do they want to build, and how do they prefer to learn?
- What concerns or blockers do they face when using AI?
- What motivates them to learn more about AI, and how likely are they to engage in AI literacy programs?

“This is a transformation moment for the entire industry, like when computing first arrived. Agents won’t replace people; they’ll make every employee a super-employee.”



Surojit Chatterjee is the founder and CEO of Ema. Previously, he guided Coinbase through a successful 2021 IPO as its Chief Product Officer and scaled Google Mobile Ads and Google Shopping into multi-billion dollar businesses as the VP and Head of Product. Surojit holds 40 US patents and has an MBA from MIT, MS in Computer Science from SUNY at Buffalo, and B. Tech from IIT Kharagpur.

Note: This report draws on insights from eight leaders who appeared on DataCamp webinars and the DataFramed podcast. Clearly indicated in the text, some quotations have been lightly edited for clarity and flow.

Defining the modern agent

There's consensus that agents blend autonomy, reasoning, and action. But not all systems claiming the title live up to it.

Others argue that the definition is widening as models evolve.

Call them reasoning systems, copilots, or co-workers, the essence is the same: they think before they act.

“[Large language models] aren't really agents. They don't know they're solving a problem, they just predict the next word.”

— Paraphrased from [The Past and Future of Language Models with Andriy Burkov, Author of The Hundred-Page Machine Learning Book](#)



Andriy Burkov is the author of *The Hundred-Page Machine Learning Book*, *The Machine Learning Engineering Book*, and *The Hundred-Page Language Models book*. His work has impacted millions of machine learning practitioners and researchers. He holds a Ph.D. in Artificial Intelligence and has successfully led dozens of production-grade AI projects in different business domains at Fujitsu and Gartner.

“An agent is anything that actively reasons, formulates a plan, executes it, and revises based on feedback.”

— Paraphrased from [RAG 2.0 and The New Era of RAG Agents with Douwe Kiela, CEO at Contextual AI & Adjunct Professor at Stanford University](#)



Douwe Kiela is the CEO and co-founder of Contextual AI, a company at the forefront of next-generation language model development. He serves as an Adjunct Professor in Symbolic Systems at Stanford University, where he contributes to advancing the theoretical and practical understanding of AI systems. Previously, Douwe was the Head of Research at Hugging Face and was a Research Scientist and Lead at Meta's FAIR (Fundamental AI Research) team.

From code to customers: Agents in the wild

Customer knowledge, scaled

In [RAG 2.0 and The New Era of RAG Agents](#), Douwe Kiela discusses Qualcomm’s engineering teams utilizing contextual agents to answer complex product questions—proof that retrieval and reasoning can thrive in high-stakes environments.

Software, simplified

Robb Wilson envisions a world where every business process becomes a “skill” callable through one universal interface.

“We don’t need 20 apps. We need one interface for all of them.”

— Paraphrased from [Designing AI Applications with Robb Wilson](#)



Robb Wilson is a pioneer in UX and AI who has led innovation initiatives for Apple, Microsoft, and Adobe. He’s the CEO and Co-Founder of OneReach.ai, the award-winning AI agent orchestration platform featured in 40+ Gartner reports in 2025.

Data, democratized

Leon Gordon shows how enterprises can align agent projects with strategy, rather than chasing novelty.

“Start with business objectives. AI is powerful, but it must prove value, not just promise it.”

— Paraphrased from [Aligning AI with Enterprise Strategy with Leon Gordon](#)



Leon Gordon is a Microsoft Data Platform MVP based in the UK and the founder of Onyx Data. Leon is an Executive Contributor to Brainz Magazine, a Thought Leader in Data Science for the Global AI Hub, chair for the Microsoft Power BI – UK community group, and the DataDNA data visualization community.

Discovery, accelerated

From Sandbox AQ’s chemistry breakthroughs to NVIDIA’s GPU-powered data-science copilots, agents are collapsing iteration cycles and unlocking creativity.

“Agents now assist the entire data-science workflow—turning ideas into experiments in minutes.”

— Paraphrased from Chris Deotte in [Data Science Trends from 2 Kaggle Grandmasters](#)



Chris Deotte is a senior data scientist at NVIDIA. He holds a Ph.D. in computational science and mathematics, with a thesis on optimizing parallel processing, and is a Kaggle 4x Grandmaster.

“Agents are turning impossible problems into merely hard ones.”

— Paraphrased from [AI Agents For Business: AI Agents and the Future of Work](#) with Marianna Bonanome



Marianna Bonanome is a leader in AI strategy and partnerships and a key member of SandboxAQ’s central AI team—spearheading the development of AI-first products that org-wide drive innovation. With over two decades of experience in mathematics, quantum research, and founding edtech companies, Marianna melds deep technical expertise with visionary leadership to craft cutting-edge AI solutions.

How agents think: The new tech stack

The AI stack is evolving from “chat plus database” to retrieval plus reasoning plus reinforcement.

- **Retrieval:** Connecting knowledge to queries in real time. “RAG 2.0” pipelines, Kiela notes, synchronize data stores, context windows, and rerankers for accuracy and compliance
- **Reasoning:** Shifting compute from training to test-time thought. Burkov highlights DeepSeek R1 and GPT-O1 models illustrate how thinking at inference boosts performance without ever-larger data sets
- **Reinforcement:** Teaching systems by reward, not repetition. “It’s the closest kind of learning to how humans and animals learn,” says Burkov on feedback-driven and self-correcting

Together, these components create agents that don’t just recall; they reason.

The human-AI partnership

Across interviews, one principle repeats: the future is cooperative.

Leon Gordon, CEO at Onyx Data, adds that adoption isn't a technical issue; it's cultural.

Organizations must invest in data literacy, transparency, and change management so employees can trust and guide their new digital teammates.

“Agents aren’t replacing humans, they’re making every employee a super-employee.”

— From [Can We Create A Universal AI Employee? with Surojit Chatterjee](#)

“Humans who use AI will replace humans who don’t.”

— From [AI Agents and the Future of Work with Sanjay Srivastava](#)



Sanjay Srivastava is the Chief Digital Strategist at Genpact. He works exclusively with Genpact’s senior client executives and technology leaders to mobilize digital transformation at the intersection of cutting-edge technology, data strategy, operating models, and process design. Previously Chief Digital Officer at Genpact, Sanjay built out the company’s offerings in AI, data and analytics, automation, and digital technology services.

“High-stakes use cases need auditability and data-aware entitlements.”

— Paraphrased from [RAG 2.0 and The New Era of RAG Agents](#) with Douwe Kiela

“When agents pass outputs to each other, their mistakes don’t cancel—they amplify.”

— Paraphrased from [The Past and Future of Language Models](#) with Andriy Burkov

The governance imperative

When agents act, accountability matters. Enterprise leaders cite three non-negotiables: accuracy, attribution, and access.

Every expert stresses the same remedy: human-in-the-loop oversight. Even autonomous systems need human judgment for ethics, edge cases, and empathy.

What's next

“Autonomous companies are closer than people think—agents running business logic, humans driving innovation.”

— Paraphrased from [Designing AI Applications with Robb Wilson](#)

Smarter thinking

Explained by Andriy Burkov and Douwe Kiela on the DataFramed podcast, reasoning-time models shift intelligence from pretraining to on-the-fly deliberation.

In other words, AI that plans, debates, and revises before answering.

Smarter senses

“Chart understanding and multimodal reasoning are the next enterprise unlock,” says Kiela. Agents will soon read slides, analyze visuals, and summarize insights across formats with ease.

Smarter systems

Wilson and Chatterjee foresee organizations where agent networks handle workflows end-to-end, freeing humans to focus on strategy, creativity, and connection.

5 leadership moves for the agentic age

The emergence of AI agents is both a technological milestone and a management challenge. Across eight thought leaders, one message was consistent: success with agents won't come from building faster models, but from building smarter organizations around them.

Five uncovered leadership moves that define the agentic age:

1

Start with pain points, not platforms

Automate what frustrates people most.

2

Design for oversight

Agents need testing and governance as much as humans need managers.

3

Invest in data fluency

Every employee is now a prompt engineer.

4

Scale responsibly

Pilot, measure, refine, then automate.

5

Celebrate co-creation

The winning cultures will treat agents as partners, not threats.

Closing thoughts

Go further

Prepare for the age of AI agents. Upskill your entire organization in data and AI skills.

[Get a Demo](#)

More of a hands-on type? Dive into creating AI agents with a code-along:

Creating AI Agents on AWS	Creating a Podcast Generation AI Multi-Agent with CrewAI
Build a Portfolio Management AI Agent with Google Gemini and Vertex	Creating a GitHub Issue Resolver AI Agent with Haystack
Building a Deep Research AI Multi-Agent with LlamaIndex	Building a News Fact-Checker AI Agent
Creating an AI Agent for Financial Report Analysis	Introduction to Creating AI Agents in Python (Part 1: Concepts)
Introduction to Creating AI Agents in Python (Part 2: Hands on)	

