CBINSIGHTS

Future of the workforce:

How AI agents will transform enterprise workflows

We dive into what agents are, their limitations, key companies, and implications for the future of work.

The idea of autonomous AI agents — LLM-powered bots that can independently reason and execute tasks — caught on like wildfire in 2023, marking an important evolution beyond chatbots and copilots.

OpenAI CEO Sam Altman has described agents as "AI's killer function" as recently as May 2024.

Now, the space is seeing a proliferation of startups: More than 50 companies have emerged since 2022 focusing on agents, agentic workflows, and agent infrastructure.

While much of the tech remains limited in its ability to execute tasks reliably, use cases are gaining traction in horizontal enterprise applications like customer support, sales, and engineering.

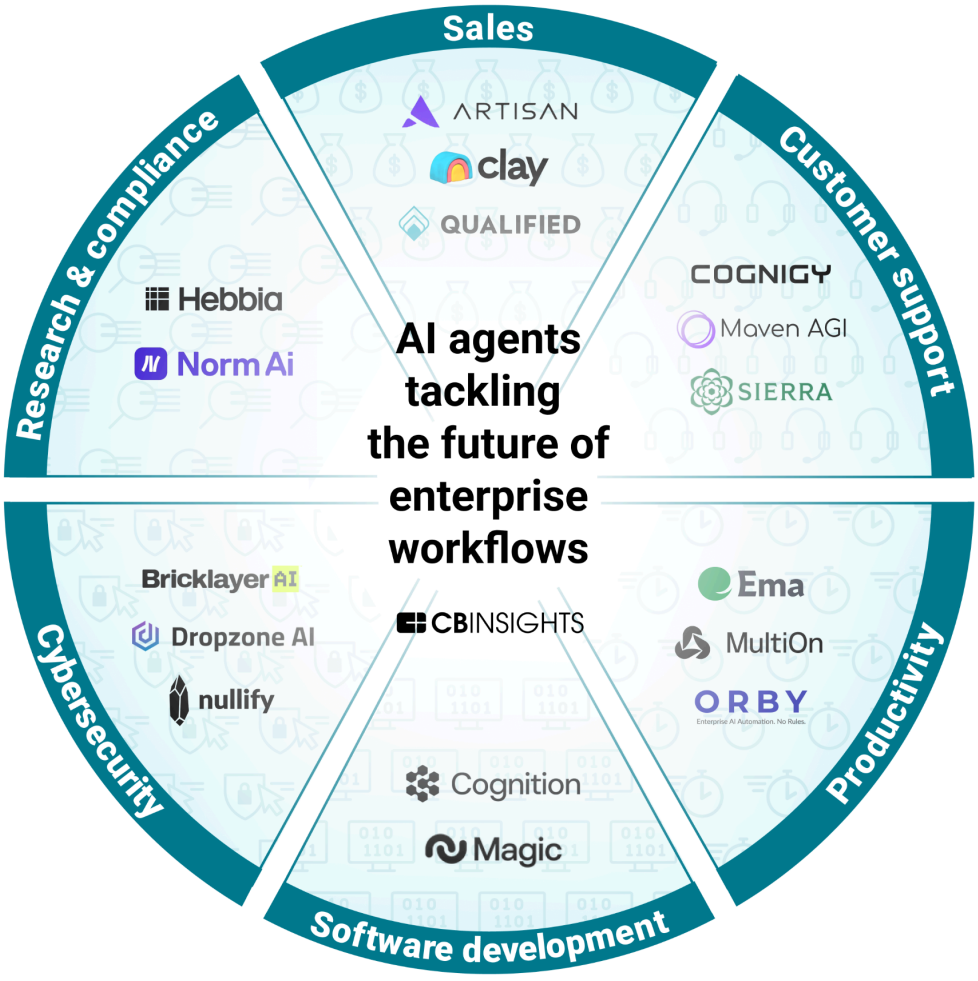

Includes leading privately held startups developing AI agents based on a minimum Mosaic score of 500. Not exhaustive of companies in the space.

CBINSIGHTS

As AI agents evolve to tackle more complex workflows, the implications for enterprises and their workforces will be far-reaching, from replacing customer service reps and SDRs to augmenting compliance efforts and more.

Below, we use CB Insights data to dig into:

- What AI agents are, current limitations, & the companies emerging in the space
- Horizontal applications gaining traction
- Industry applications on the horizon
- Looking ahead

Contents

Note: Companies profiled in this report are working on various levels of autonomy, from agentic, LLM-powered workflows to fully autonomous agents.

The state of AI agents

Generative AI has revolutionized language processing, enabling tasks from creative writing to code generation, with ChatGPT's 2022 release bringing the tech to hundreds of millions of users.

Since then, a range of chatbots and copilots have emerged, leveraging LLMs to answer user questions, summarize documents, complete software code, and more.

Agents mark an important evolution beyond chatbots and copilots: AI agents can tackle complex tasks on a user's behalf. They can use websites, schedule meetings, plan vacations, and more.

While AI agents haven't fully taken off since interest in them surged in 2023, the tech is evolving quickly and becoming more capable — with varying degrees of autonomy. We examine what agents are, their limitations, and the landscape below.

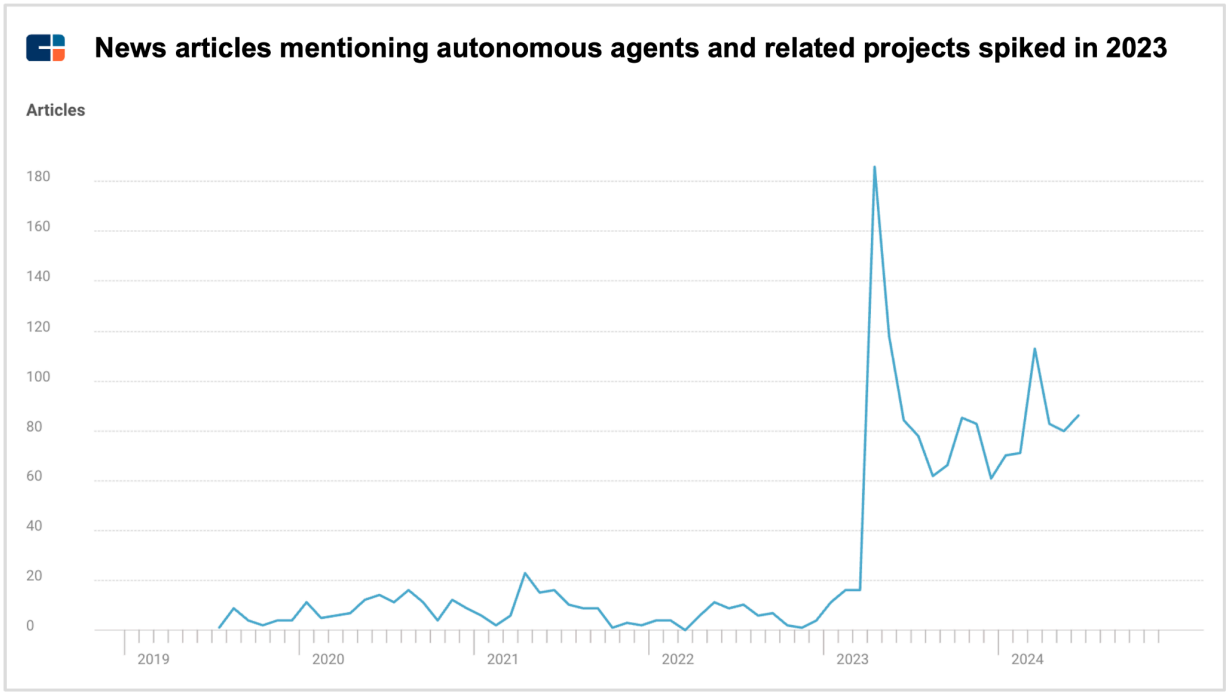

Source: CB Insights [news mentions search]

How AI agents work

Autonomous agents are AI systems that can pursue a goal without needing humans to tell them what to do and when. If given an objective, they can create tasks, complete those tasks, and re-prioritize them until reaching the objective.

Integral components of these systems include foundation models (for reasoning and analysis), tool use (to interact with the internet and other software or apps), and memory access.

Early projects released in 2023 highlight this framework, from the open-sourced AutoGPT, which saw massive interest from the developer community, to its simpler cousin BabyAGI.

For example, BabyAGI (illustrated below), when given an objective by the user, will loop to execute on it, relying on "agents" (e.g., OpenAI's GPT-4 model) to complete tasks, pulling in context from the LLM and storing completed tasks in memory (using a vector database).

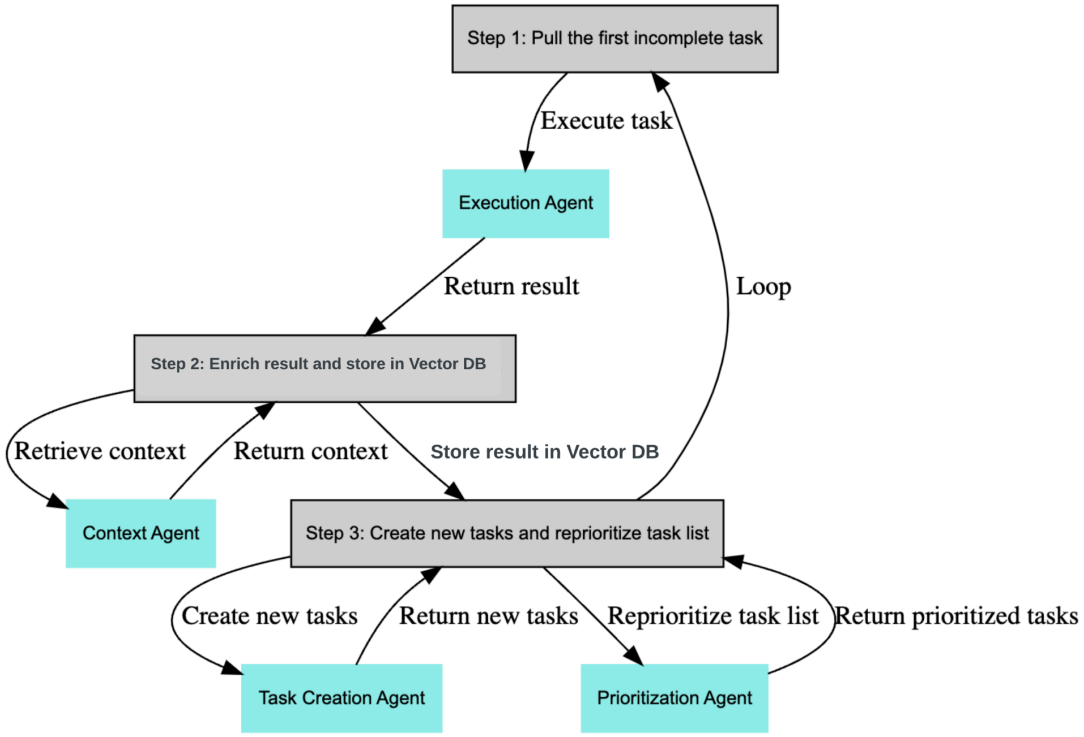

Source: Yoheinakajima/BabyAGI

Limitations include complex reasoning & planning and tool use

Despite compelling demos — from booking flights to conducting market research — fully autonomous agents are still in the early adoption phase.

Current agent projects are limited in their ability to execute tasks reliably, especially those aiming to work across the entire internet with its array of websites, error states, and unpredictable user interfaces.

One of the most advanced software engineer agents — Cognition's Devin — is able to resolve just 14% of issues unassisted, based on SWE-bench, a benchmark that evaluates agents' ability to resolve real GitHub issues in open-source projects. This still places it far ahead of other models.

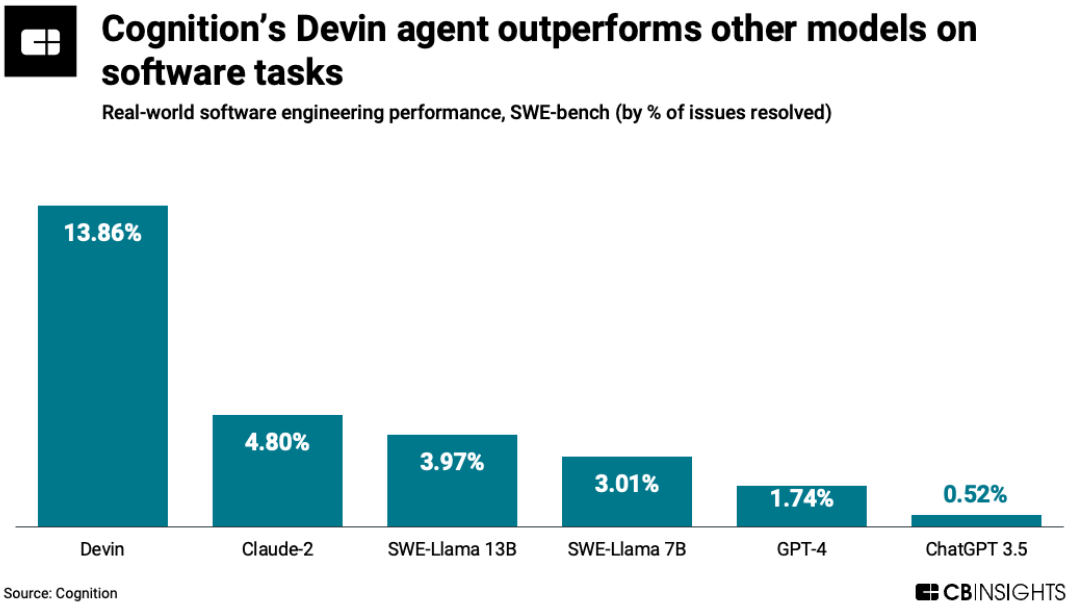

Source: Cognition

Meanwhile, OpenAI reportedly still considers its technology to be at the "chatbot" stage (AI with conversational language), though it claims it is approaching "reasoners" (problem-solving like a human). The company reckons that the third level — full-fledged agents that can handle complex actions for users — has not arrived yet.

One challenge is getting LLMs to interact accurately with webpages and outside tools through APIs — a key feature of agents expected to perform complex tasks across software platforms. They also need to chain multiple functions logically, with no room for hallucination.

For example, AI devices like Rabbit's R1 and the Humane Ai Pin, which rely on large models, have struggled to actually follow through on tasks, like ordering an Uber or queuing up songs on Spotify when asked by the user.

However, some startups are exploring workarounds, such as optimizing agents to handle only a limited set of tasks, keeping humans in the loop to give feedback and steer the agent's approach, or deploying a secondary agent to check the first one's work.

For instance, Sierra, which is focused on customer support applications, uses a multi-agent approach. One Sierra customer we interviewed described it as a "check system":

The other thing that's really unique about them is **they've got a check system. They don't have just one AI agent working on our project. They have two.** Before it actually responds to you, in a nanosecond, it bounces the answer off of another AI agent. That agent says, "Yes, this sounds like something that [the brand] would say," and then it feeds it back and then it responds to you. **There's this checks and balances system that allows it to make sure that it's actually responding appropriately to those answers.** If it's not, it'll go back and work out the problem with itself, and then it'll go back to the consumer, or it'll transfer it to a human agent if it doesn't feel like it can adequately answer it.

Director of Marketing,
Consumer goods company

The AI agent landscape

Horizontal applications dominate

Alongside the active community of open-source AI agent projects, the landscape of VC-backed agent startups is dominated by a focus on horizontal applications — across sales, customer support, and other enterprise and general productivity workflows.

Startups focusing on vertical-specific applications remain limited today based on our analysis.

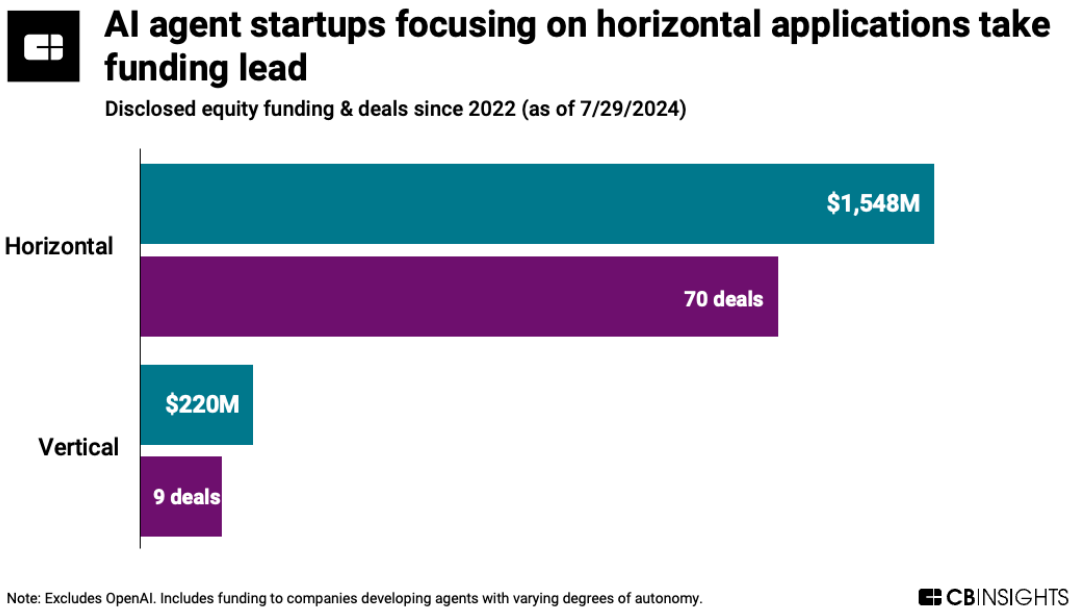

The vast majority of startups developing agents are targeting businesses, not consumers. They're vying to replace the roles of SDRs, software developers, compliance analysts, customer service reps, and more.

The implications will be far-reaching — from augmenting or replacing companies' employees day-to-day to eventually becoming the outside face of companies to their customers.

Agent infrastructure tools emerge

Companies in the AI infrastructure space are making it easier for developers to build AI agents.

For example, LangChain offers an open-source framework for building LLM applications, including agents. Meanwhile, Y Combinator-backed Zep AI, which integrates with LangChain, is building "long-term memory" for developers' agent applications.

Emergence AI, a division of Merlyn Mind that came out of stealth in June 2024 with nearly $100M in funding, is building an "Orchestrator agent" that routes tasks to the best LLM or agent.

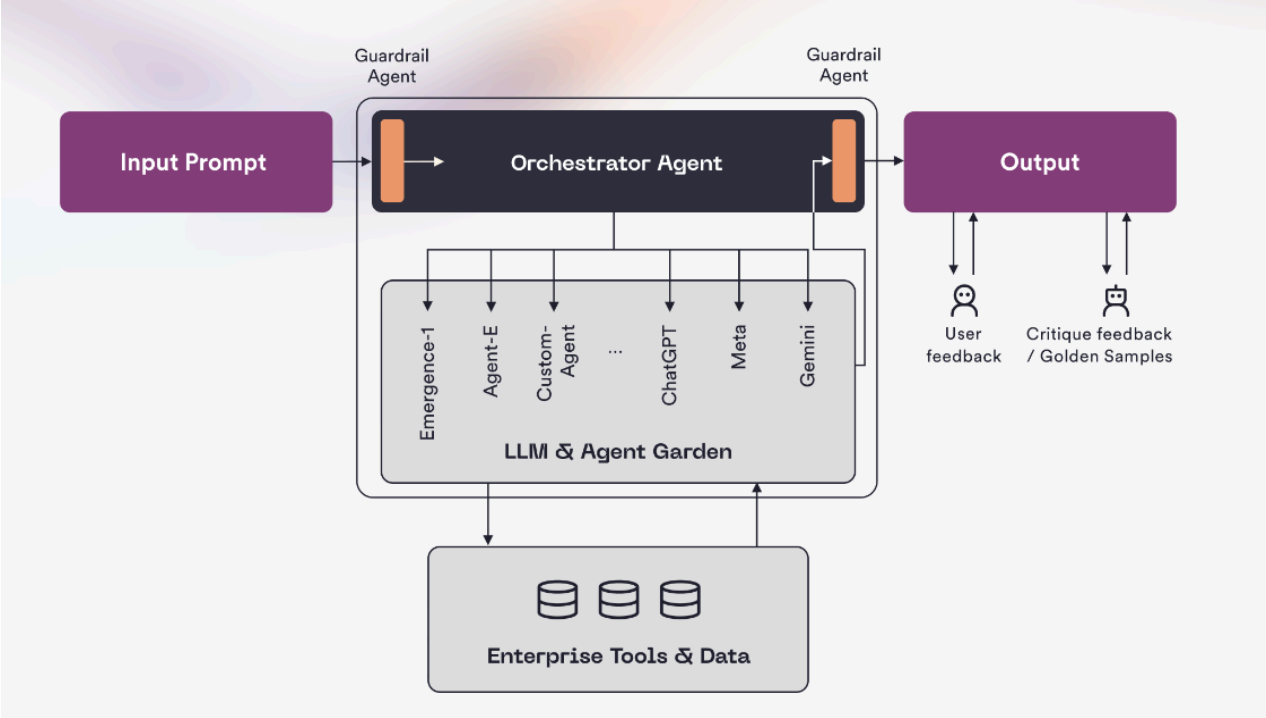

Source: Emergence

Startups are also emerging to make it easier for AI agents to interact with the web.

For example, Browserbase's platform allows developers to automate web interactions with AI and headless browsers (web browsers without a graphical user interface that can be controlled programmatically), as highlighted by its CB Insights Scouting Report.

Browserbase's platform enables automated web interactions with AI and headless browsers

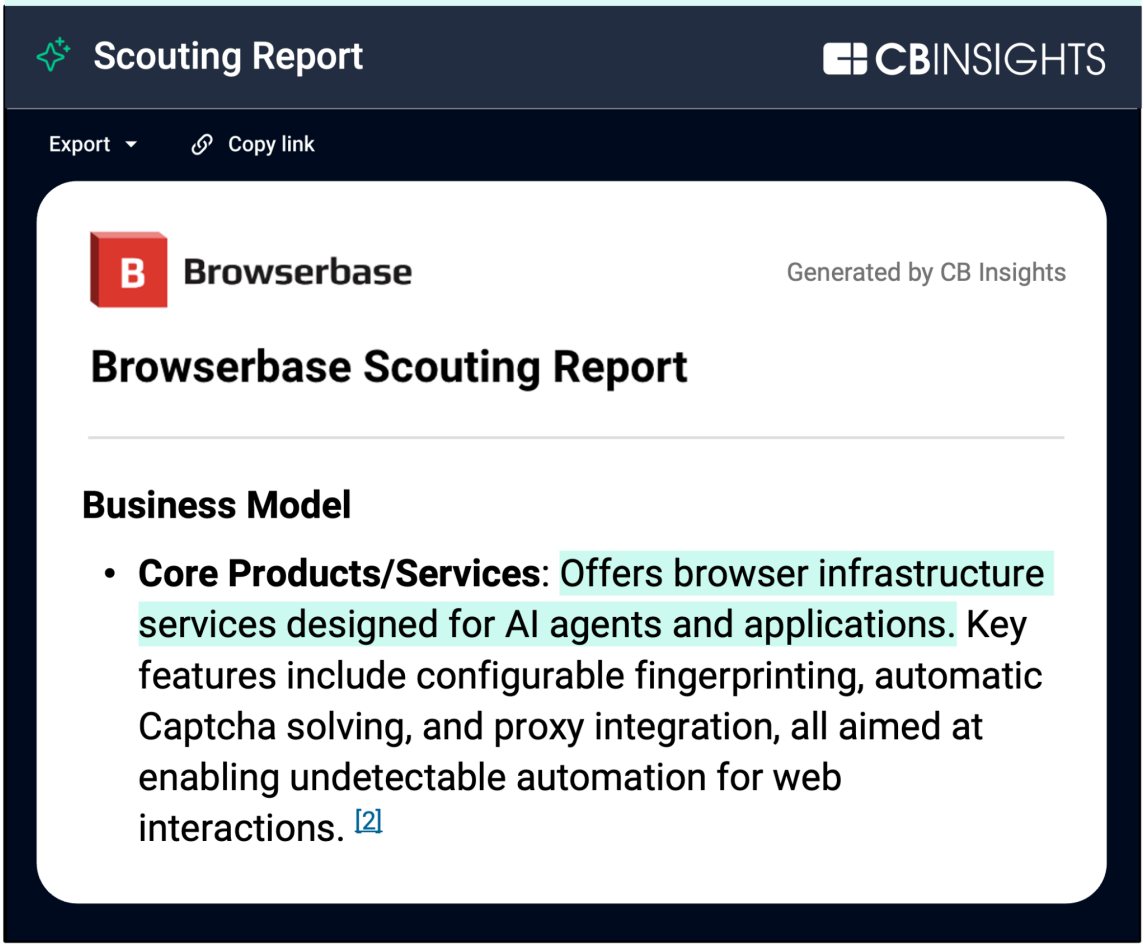

Source: CB Insights Browserbase Scouting Report

Another company here, Anon, was founded to help solve authentication challenges that AI agents face online, such as entering usernames and passwords securely. The company offers developers a software development kit (SDK) that enables agents to authenticate on behalf of their users as they navigate websites.

Investment surges

Companies focusing on agents are seeing surging levels of investment and deals in 2024 so far (this excludes OpenAI, which works across numerous AI areas).

Top-funded companies include Adept ($415M in total funding), Imbue ($232M), and Cognition AI ($196M).

Notably, Adept's leadership was effectively acqui-hired by Amazon in June 2024, though the remaining team plans to focus on agentic AI.

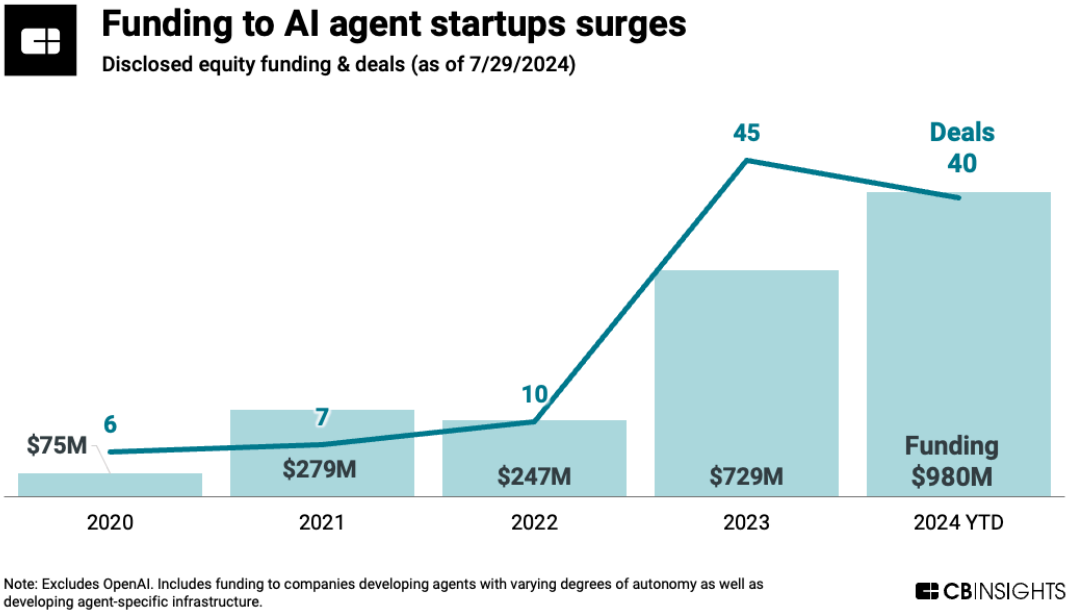

Source: CB Insights advanced search

The most active investors in the space are Y Combinator, having backed 12 startups here, and Abstract Ventures, with 5.

Corporate investors are also eyeing the space: Citi Ventures and New York Life Ventures both backed compliance-focused Norm Ai, while Workday Ventures and Atlassian Ventures invested in Adept.

Big tech jumps into the fray

Microsoft, Google, and Amazon all have agent projects in motion.

- **Microsoft** announced in May 2024 an extension of its Copilot tools as "team members," with new capabilities to develop agents that automate workflows.
- **Google's** DeepMind debuted Project Astra at its developer conference in May 2024. Project Astra is a "universal AI agent" that functions as a multimodal AI assistant, able to interact with users via video, image, speech, and text. Alphabet CEO Sundar Pichai described the project on the company's Q2'24 earnings call as a "glimpse of the future":

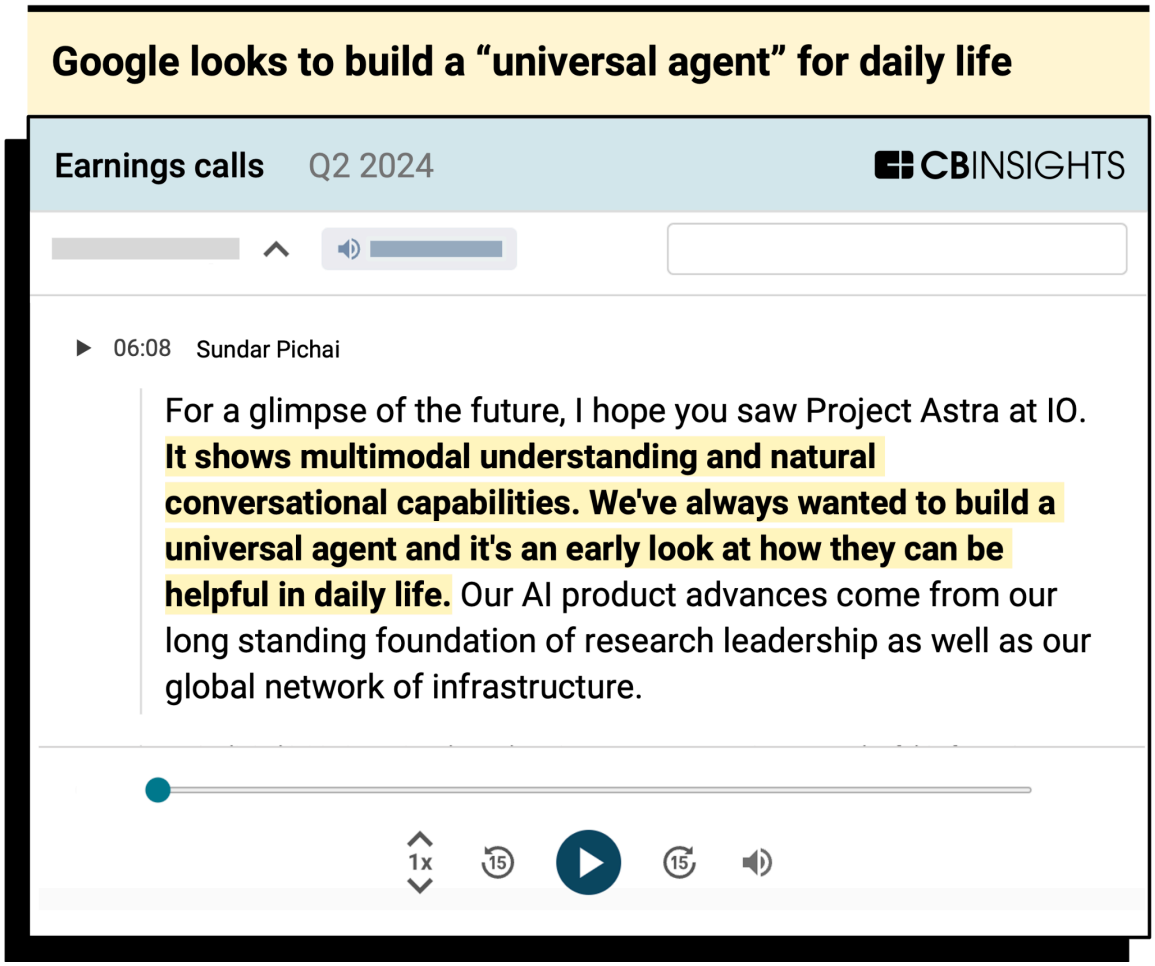

Source: CB Insights [Alphabet earnings call transcripts]

- **Amazon** brought on Adept’s co-founders and will license its tech as of June 2024. In a memo to employees, Rohit Prasad, who heads up Amazon’s Artificial General Intelligence (AGI) team, said the move “will accelerate our roadmap for building digital agents that can automate software workflows.”

On the investment side, Google Ventures has backed Hebbia and Cognosys. The Amazon Alexa Fund also invested in MultiOn’s January 2024 seed round and Imbue’s Series B.

Horizontal applications & impacts

The vast majority of AI agent startups are focused on automating workflows that cross company departments. Leading categories, which we dive into below, include customer support, sales, and software development.

Customer support reps are at high risk of AI automation

AI has been a feature of the contact center and customer support landscape for years.

For example, one buyer we spoke with replaced 135 full-time employees (FTE) involved in call routing in 2019 by leveraging Google's Contact Center AI Dialogflow solution.

Google Cloud

> The Google Dialogflow solution was part of our voicebot, and what was driving the voicebot was a broad, simple use case of switchboards. We would take calls into the stores, and then the switchboard would understand the caller intent and route the caller directly to the extension, managing the caller being serviced at that particular extension. **There were around 135 FTE in this activity, and we replaced them with a chatbot. As part of the chatbot, we needed an inferencing engine or machine learning/AI to understand the utterance and match it to a particular intent that we could then match to a routing point to service that specific intent.**

Manager,
British multinational merchandiser

CBINSIGHTS

Source: CB Insights Google Cloud buyer interview

However, chatbots have historically been limited by rigid, rule-based responses and a lack of contextual understanding.

Now, LLMs are enabling much deeper automation and cost savings by handling more complex conversations and workflows, in turn freeing up agents' time (or replacing them altogether).

In February 2024, Klarna announced its OpenAI-powered AI assistant was doing the work of 700 full-time customer service agents after being live for 1 month. The assistant is expected to deliver $40M in annual savings.

In our conversations with customers, it's clear that AI agents will increasingly handle the tasks of human support agents — a role in which nearly 3M Americans are employed, as of 2022, per Bureau of Labor Statistics (BLS) data.

Sierra, co-founded by former Salesforce CEO Bret Taylor, has raised $110M since its founding in 2023 to build customer service AI agents. This Sierra retailer customer is leveraging the company's AI agents to deal with the "need to flex temporary teams during high volume periods":

SIERRA

My stretch goal for everybody is 4.75/5 CSAT. We're consistently seeing the AI agents and our human agents deliver around a 4.5. It's competing and it's on par with what our human agents are doing…The second KPI we're looking at is actual cost per ticket handled…**I want to make sure that the cost of Sierra is equal to or lower than the human agent cost**. We're also looking at resolution rate — the percentage of tickets that the AI agent is getting and resolving itself 100% without needing to escalate to a human agent. **Right now, we're seeing about a 60% resolution rate, which is great because it's made me not have to hire the temporary staff** on the ticket side that I've hired in previous summers.

Director of Marketing,
Consumer goods company

CBINSIGHTS

Source: CB Insights Sierra buyer interview

An older group of AI customer support companies, which emerged prior to LLMs' mainstream arrival, are also pivoting to focus on AI agents. These include Ada, Forethought, and Intercom, as well as Cognigy, which raised $100M in Series C funding in June 2024 to scale its contact center agents.

Looking ahead, expect voice solutions to come more into focus. For instance, Fixie, which initially developed chat-based AI agents, has switched its focus to real-time voice interactions.

Meanwhile, this Sierra buyer said,

> *"**I expect our AI usage to really increase as Sierra onboards voice**, which they're thinking about doing by the end of [2024]. I'm considering not turning it on full time, but potentially turning it on and off hours when our agents aren't working to be able to solve those consumer inquiries when we don't have active phone support available."*

Other buyers are looking to leverage AI agents to bring new levels of personalization to their customers outside of just conversational support interactions:

> *"How can we leverage AI to predict preferences or lead to tailored product recommendations or personalization in marketing, etc.? I see that's one potential area of opportunity that could be a part of a customer experience Sierra model."*

Sales "AI SDRs" is a crowded market

One of the most active categories in terms of company activity is startups developing autonomous sales reps that can scale outbound automation. They research, route, and qualify leads, personalize messages, and book meetings.

Clay ($62M in total funding), for example, is developing "Claygent," its AI tool focusing on data enrichment and sales prospecting. Other established companies focusing on broader SDR workflows include Qualified ($163M in total funding), which launched Piper, its AI SDR, in April 2024.

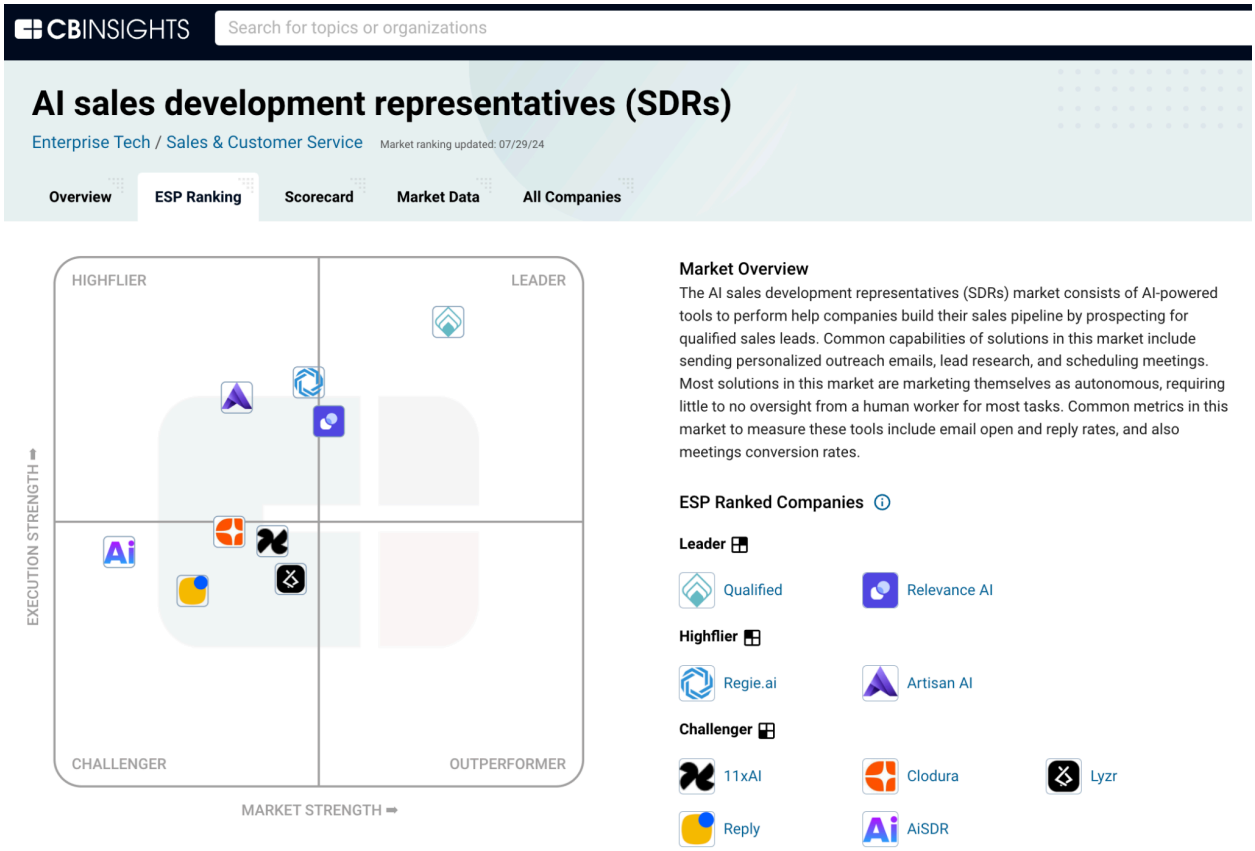

Source: CB Insights [AI SDRs market report]

Founded just last year, [11xAI] has scaled to $2M in ARR as of March 2024 with Alice, its AI SDR. The company primarily offers a task-based pricing model — charging for completed tasks like identifying and researching accounts, preparing outreach, and scheduling meetings.

For example, [one Sales Operation Manager who's a customer of 11xAI] pointed to the usefulness of the tool to replace SDRs, but also noted concerns about AI-generated email fatigue:

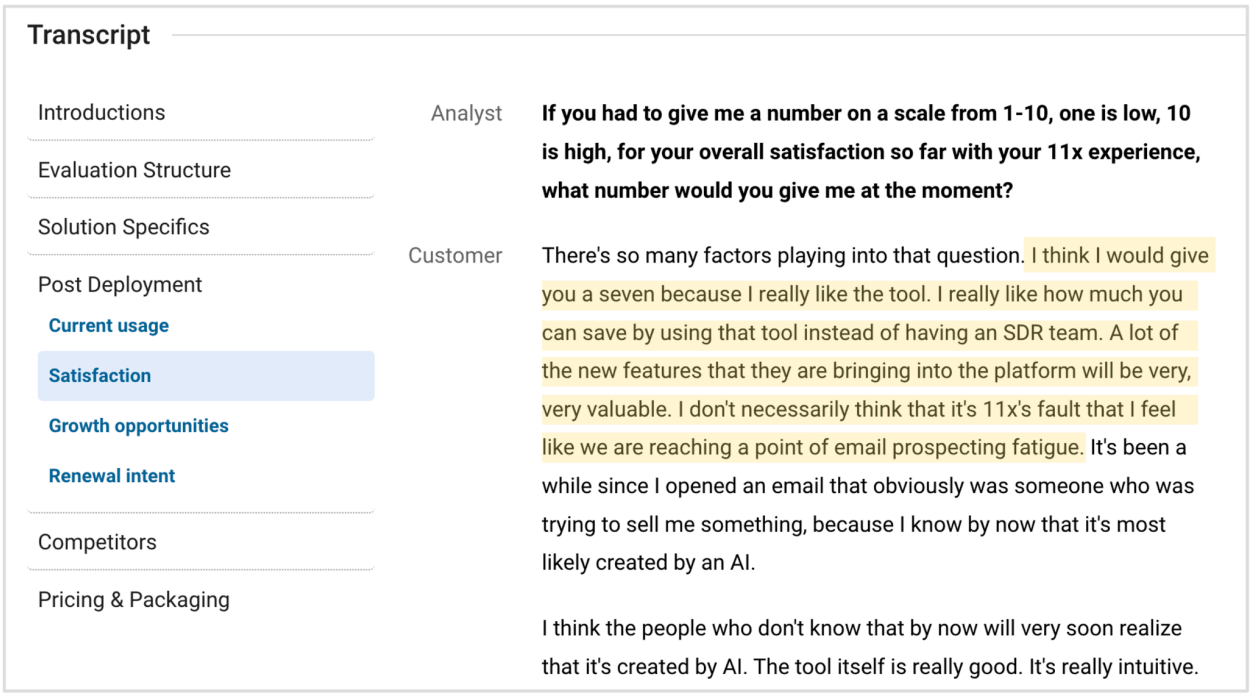

Source: CB Insights 11xAI buyer interview

Another common point of concern is the trustworthiness and quality of the messaging:

- Clay customer (CSAT: 7/10): "Another way is that when we're using the Claygent capabilities, all the different AI capabilities that you can use, even though it shows its work and it explains the logic or the rationale that it's using, a lot of times, **we just aren't fully trustworthy of it yet. We have humans double-checking every single thing.** It does the work for you, but we don't trust it yet."
- 11xAI customer (CSAT: 7/10): "**I think we lost a little faith in the AI's ability to create content that we, with peace in mind, could send out to potential prospects, which I think has taken a little bit of the groundbreaking smartness out of the tool** for us because that actually means that we would just be able to set up campaigns in any other cadence tools and run them through them."
- 11xAI customer (CSAT: 5/10): "We were also selling more to the lower, middle markets or lower companies, 100 to 150 employees, which again feels like we would get more return. **The challenge — why I'm giving it a 5 — is we are actually shifting to work on enterprise companies and thereby the messaging, the**

complexity of messaging is a lot more sophisticated. Therefore again, the returns aren't there compared to if you would message the lower / mid-market."

The prospect of automating many of the repetitive tasks SDRs do every day is driving activity and interest in the market. However, more sophisticated messaging and personalization, and better incorporation of intent signals, will be key for buyers evaluating solutions in the market.

Software development agents look to go beyond copilots

Coding copilots are becoming standard-issue tools at enterprises.

Oil & gas major BP said on a recent earnings call that it needs "70% less coders from 3rd parties" thanks to AI, as seen on the CBI platform below.

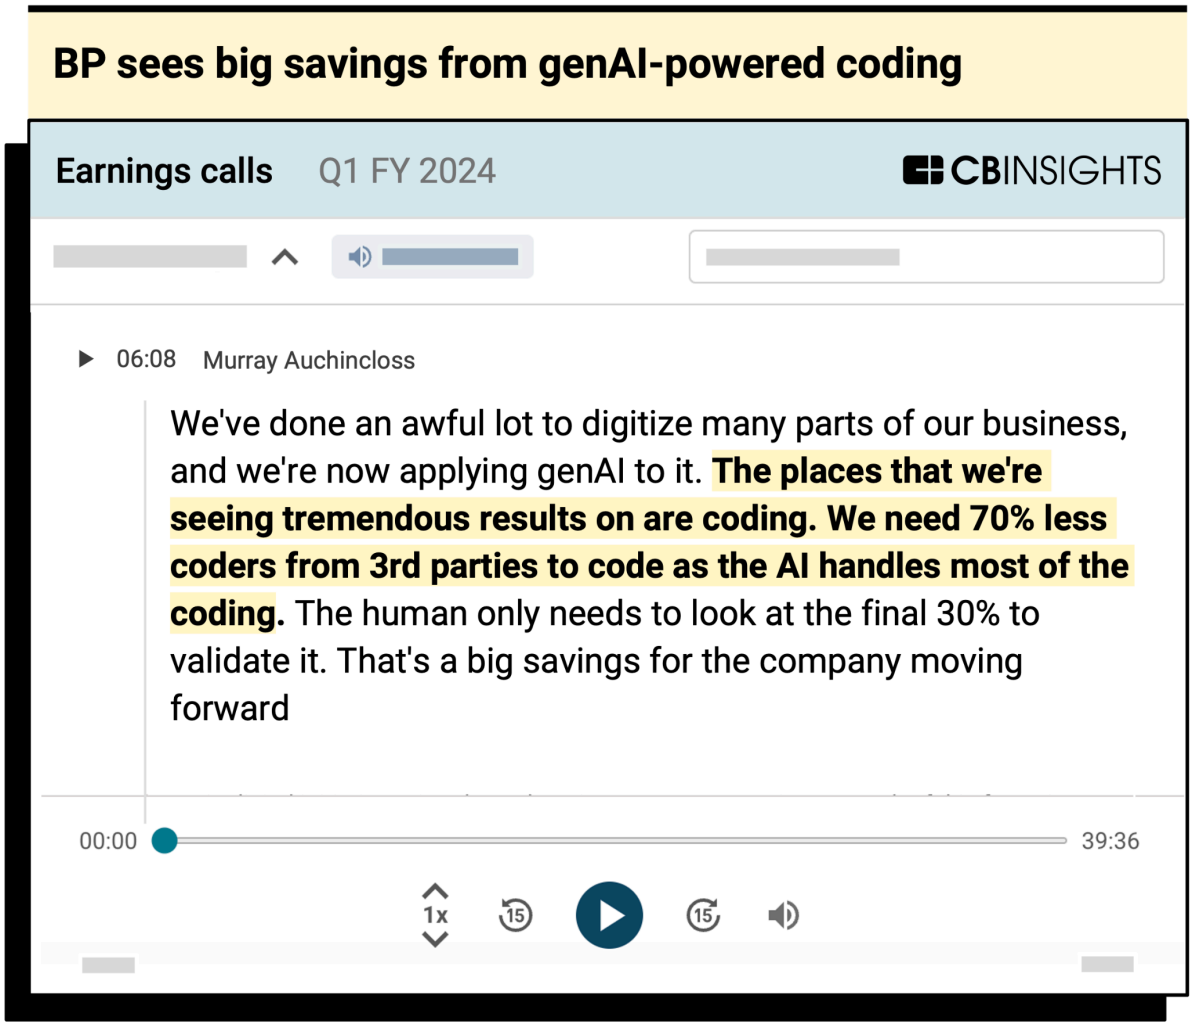

Source: CB Insights BP earnings call transcripts

The AI-powered coding assistant scene has exploded, with startups raising massive funding, alongside products from big tech like Microsoft's GitHub Copilot, Meta's open-source Code Llama, and Amazon's CodeWhisperer.

Companies like Cognition and Magic are aiming for true coding agents that can replace the work of human engineers — as opposed to assistants that suggest code snippets.

The tech is still early, and quality remains a major challenge.

This Magic customer — an AI cloud software developer lead at a Fortune 500 company — highlighted 3 issues in a recent interview: 1) lack of originality in the written code; 2) privacy and data concerns; and 3) excess "fat" in the code.

But the opportunity is large. There are 1.8M software engineers in the US with a median pay of $130K per year, per BLS data, driving a combined annual spend of over $230B. Given the high cost and extensive need for developers, we anticipate continued investment and adoption in the space. Expect other companies in the copilot space to begin developing agentic workflows as well.

Cybersecurity is an emerging category

One emerging horizontal application of agents is in security operations — for example, investigating alerts, generating reports, and eventually deploying fixes for vulnerabilities.

Nullify is building an "AI security engineer," while Dropzone AI is focused on creating AI SOC (security operations center) analysts, as highlighted by CBI's Funding Insights for its recent Series A.

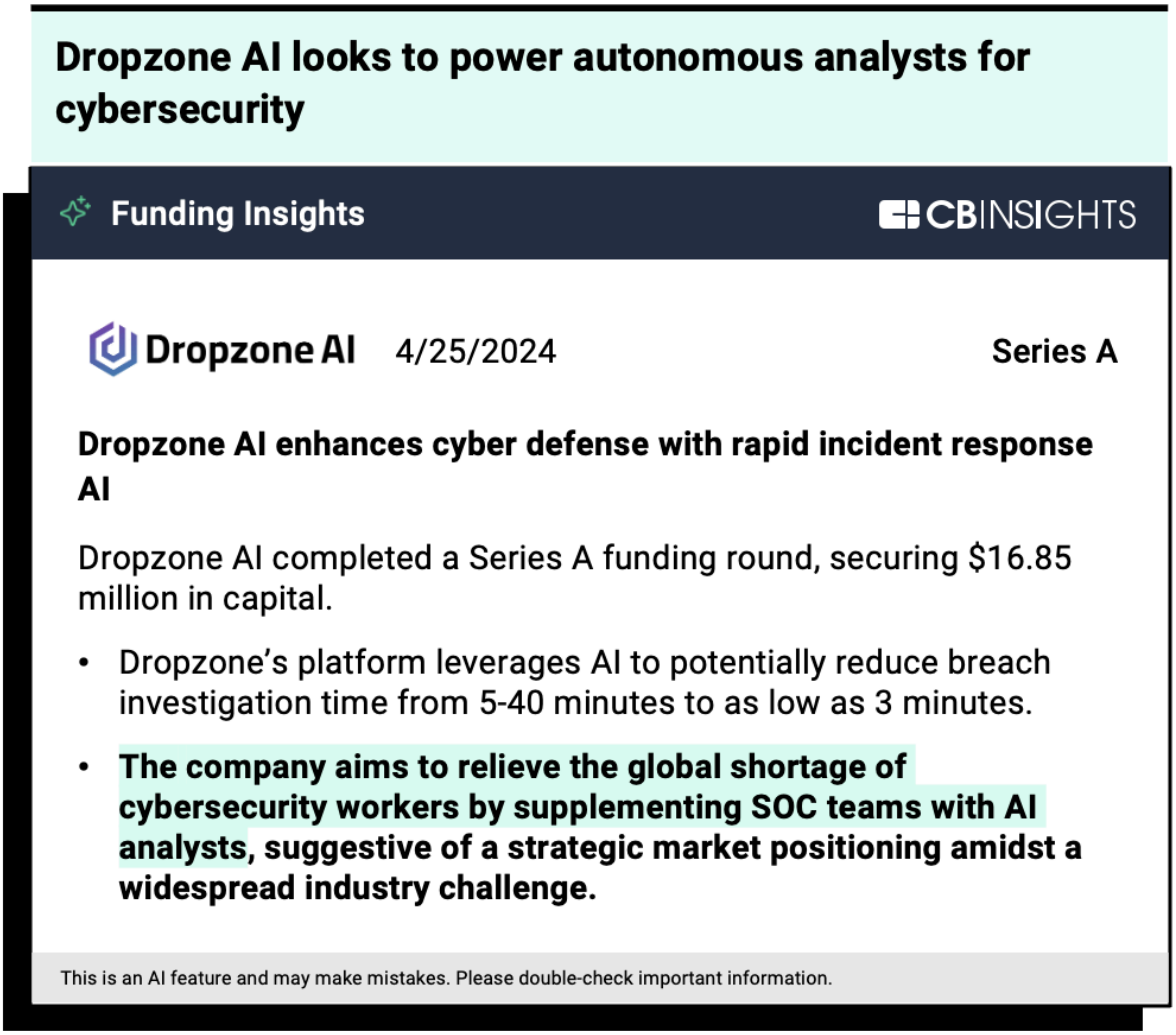

Source: CB Insights [Dropzone AI Funding Insights]

General enterprise workflow automation

One of the most robust categories of agent companies is those targeting a wide range of enterprise workflows. Applications span general productivity and research use cases.

Many startups in the category also highlight applications across functions like sales, customer service, and HR.

Select VC-backed enterprise workflow AI agents

Companies	Country	Founded Year	Commercial Maturity	Mosaic (Over...
Adept	United States	2022	3: Deploying	826
Orby AI	United States	2022	2: Validating	734
Ema Unlimited	United States	2023	2: Validating	762
Relevance AI	Australia	2020	4: Scaling	552
Respell	United States	2022	2: Validating	446
Twin	France	2023	1: Emerging	532
CognosysAI	Canada	2023	3: Deploying	419
Perhaps	Chile	2023	3: Deploying	480

Source: CB Insights advanced search

Enterprise adoption is already starting to happen, based on company customer data and our buyer interviews.

For example, Accel-backed Ema, a 2024 CB Insights AI 100 winner, advertises itself as a "universal AI employee" targeting a range of functions and industries. The company's customers include TrueLayer and Moneyview.

Another company here is early-stage Respell, which helps businesses build workflows using models from OpenAI, Cohere, Anthropic, and other providers. One Respell customer we spoke with — a director of data engineering at a $50M+ funded data platform — leverages Respell for a number of sales and marketing use cases:

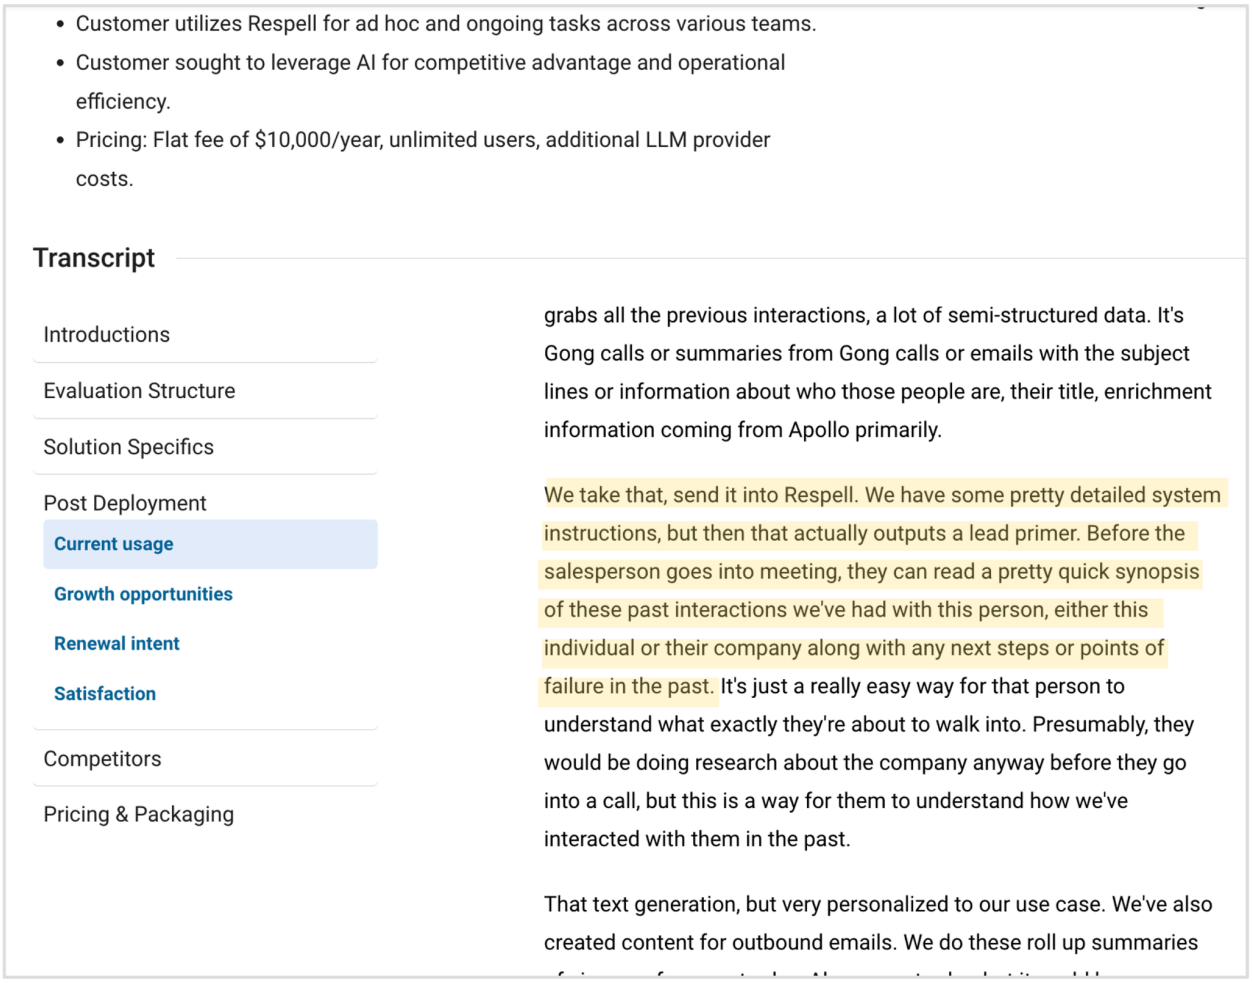

Source: CB Insights Respell buyer interview

A number of other early-stage companies in the space are working more specifically on web research and data extraction, including Zeta Labs, Reworkd, and Lutra AI.

Emerging industry applications & opportunities

While few agentic companies focus on single industries, we expect to see more companies "niching down" in the coming year, mirroring trends in the broader generative AI space.

Agents' reliability is a bigger hurdle for sensitive, highly regulated industries like financial services and healthcare. Fully autonomous workflows may be years away, but corporations should monitor developments among companies in the space.

Below, we highlight emerging applications or opportunity areas across 5 industries.

Financial services

While still early, two primary applications of agents in financial services are on the horizon: 1) compliance and 2) investment research.

In compliance, Norm Ai — whose backers include notable finserv CVCs and incumbents Citi Ventures, New York Life Ventures, and TIAA — is building agents trained on regulatory filings to make compliance determinations.

Another company here, Parcha AI, which was founded by former Coinbase and Brex employees, emerged to build AI agents focused on compliance workflows in fintech and banking.

The company explained in a blog post the challenges of building "reliable agentic behavior" for something like compliance workflow automation:

> *"While the concept of agentic behavior was promising, building reliable agentic behavior with large language models (LLMs) was a massive endeavor.* ***Creating general-purpose "autonomous agents" could have taken us years.*** *…* ***[Our customers] would much rather have a solution that was very accurate and reliable for a subset of tasks than a fully autonomous solution that could automate a workflow end-to-end but worked only 80% of the time.****"*

Parcha has since pivoted to home in on LLM-powered workflows for KYC/KYB and due diligence — rather than having a fully autonomous agent determine and take actions.

With the reliability question still unsolved, more limited workflows will be in focus. That includes investment research.

Hebbia, which raised a $130M Series B led by a16z in May 2024, is building agents for financial services. For example, its Matrix product builds spreadsheets that mine information from files (in rows) and deliver answers to questions (in columns).

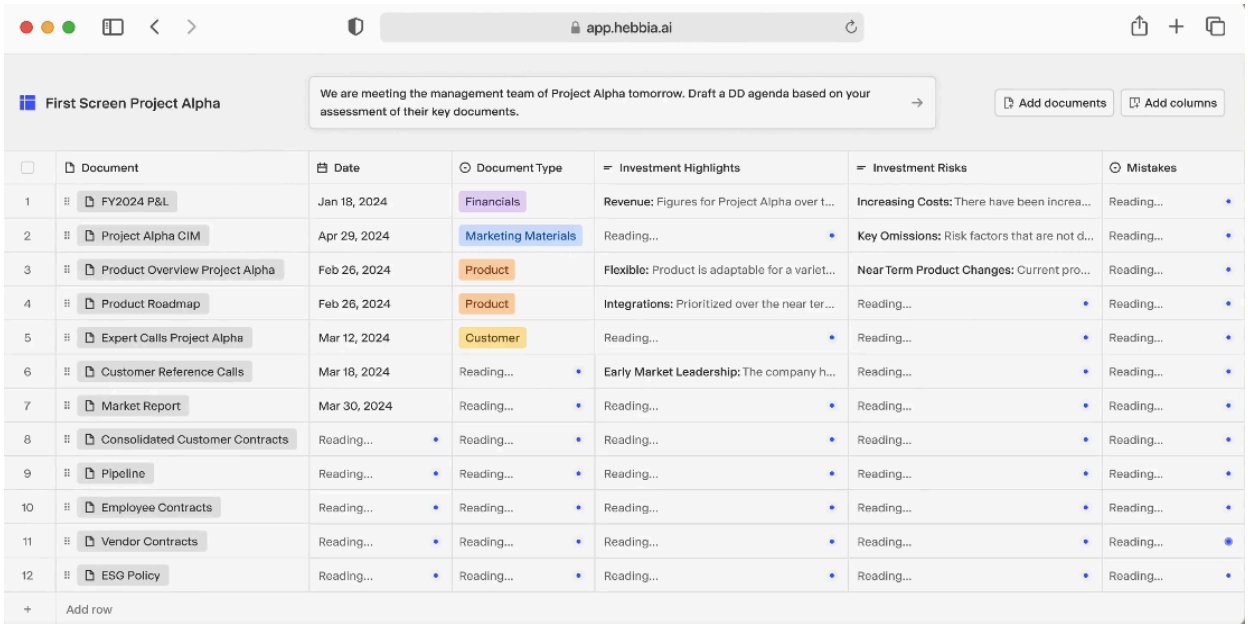

Source: Hebbia via a16z

Insurance

Applications of agents in insurance include underwriting and claims processing.

Tech providers are already eyeing genAI-powered underwriting workflows.

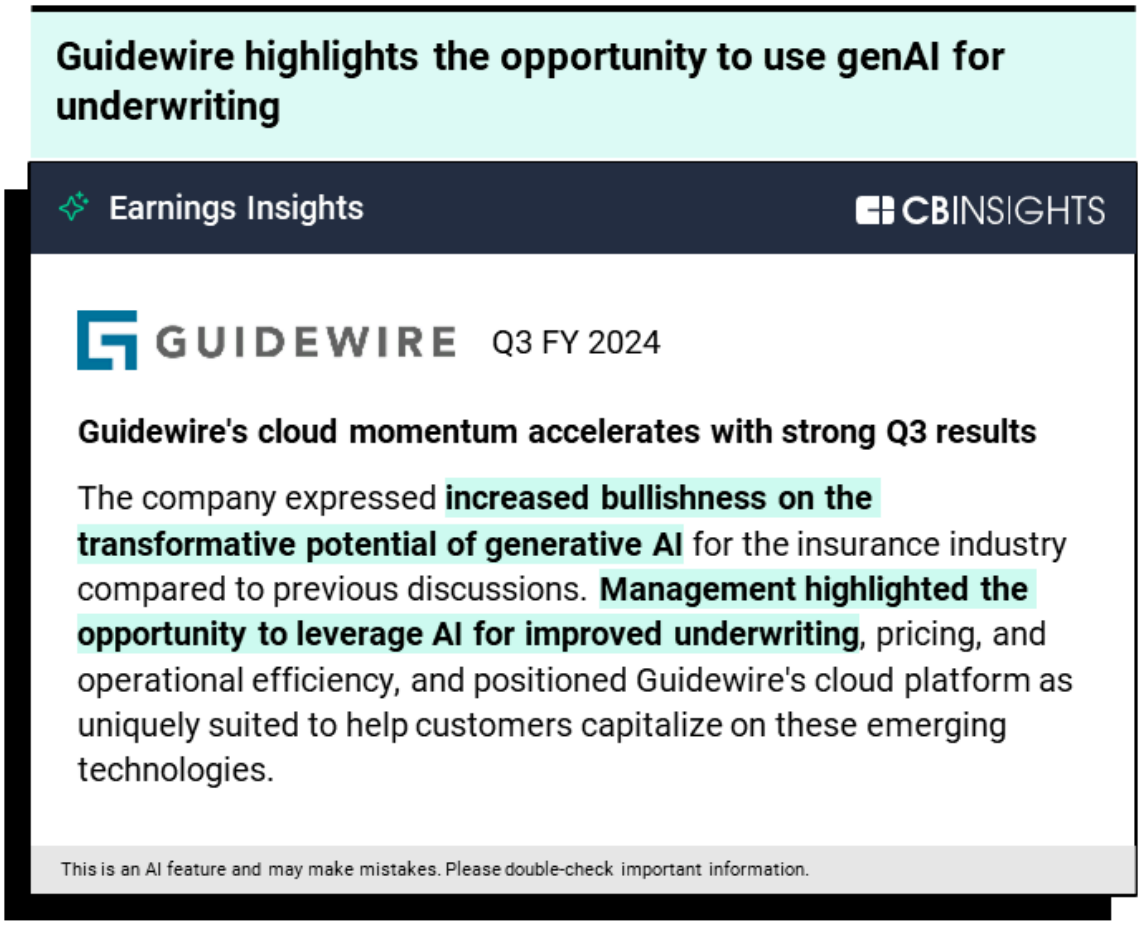

Source: CB Insights Guidewire Earnings Insights

The next step will be more agentic workflows. For example, Roots Automation — which offers an insurance-focused LLM centered on unstructured data and an AI underwriting assistant — raised a funding round from Erie Strategic Ventures in January 2024.

Healthcare

Based on our analysis of the market, autonomous solutions in healthcare have yet to emerge — though some agent companies like Ema advertise healthcare-specific applications such as a pharmacy assistant.

Monitor the medical scribes market for the development of more advanced workflows — an already active generative AI market where copilots are automating tedious tasks like documentation.

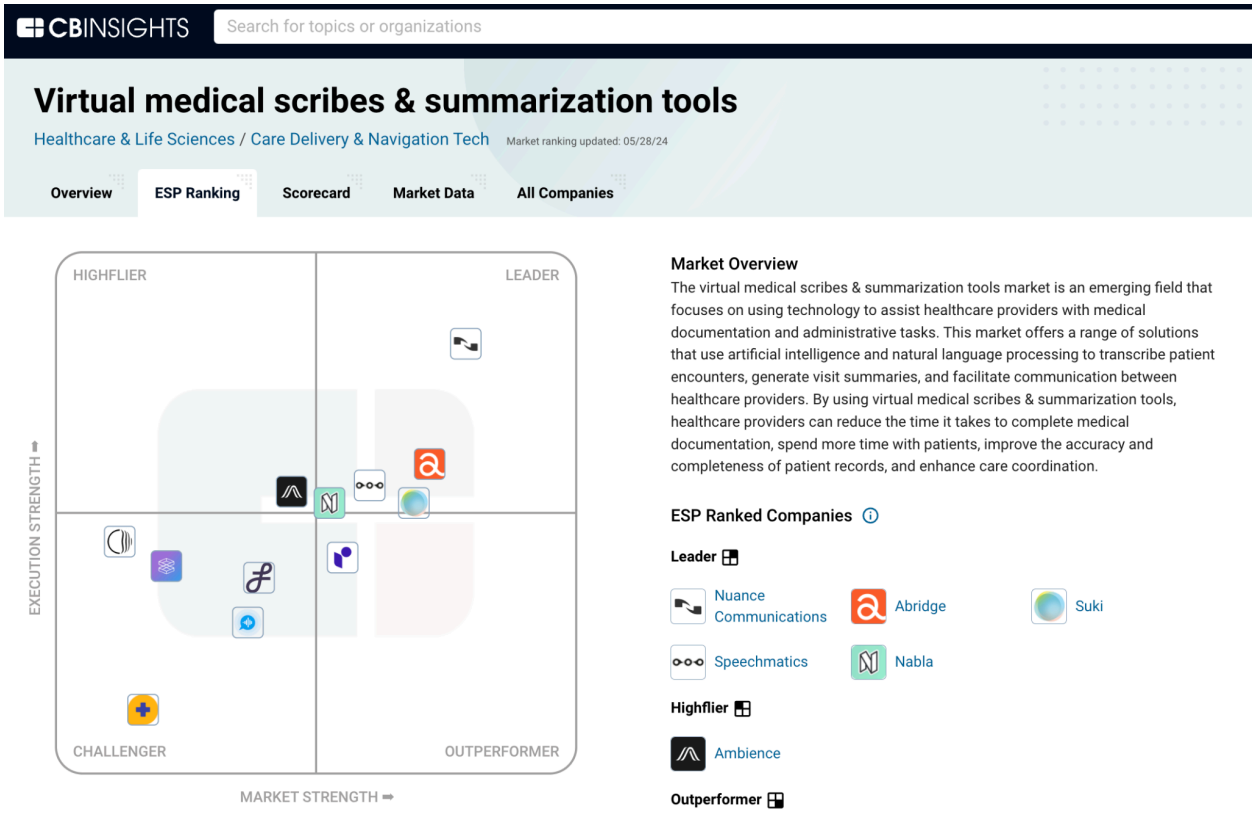

Source: CB Insights [Virtual medical scribes & summarization tools market]

Industrials

Companies are developing [industrial autonomous agents] that optimize processes and equipment — including control systems, robots, and other industrial machines — without needing humans.

For example, [Composabl] launched its autonomous agent in May 2024. It uses LLMs to automatically control and optimize industrial equipment, with early adopters like Rockwell Automation and RoviSys already experiencing success.

In the long term, as humanoid robot development progresses, humanoids are likely to gain even more advanced decision-making capabilities (i.e., agentic behaviors) to overcome unexpected challenges or adapt to their environments in ways that dramatically increase their autonomy and effectiveness.

Learn more [about applications of humanoid robots here].

Gaming

While a number of companies are developing genAI-powered NPCs (non-player characters) for games, some are going a step further.

DeepMind published research in March 2024 on a "generalist AI agent for 3D virtual settings." The agent can understand and navigate different game environments as well as carry out tasks as a human could.

Meanwhile, Altera is a research lab working on "digital human beings." Its first product is an agent that can play alongside users in Minecraft. The company raised $9M in seed funding in April 2024.

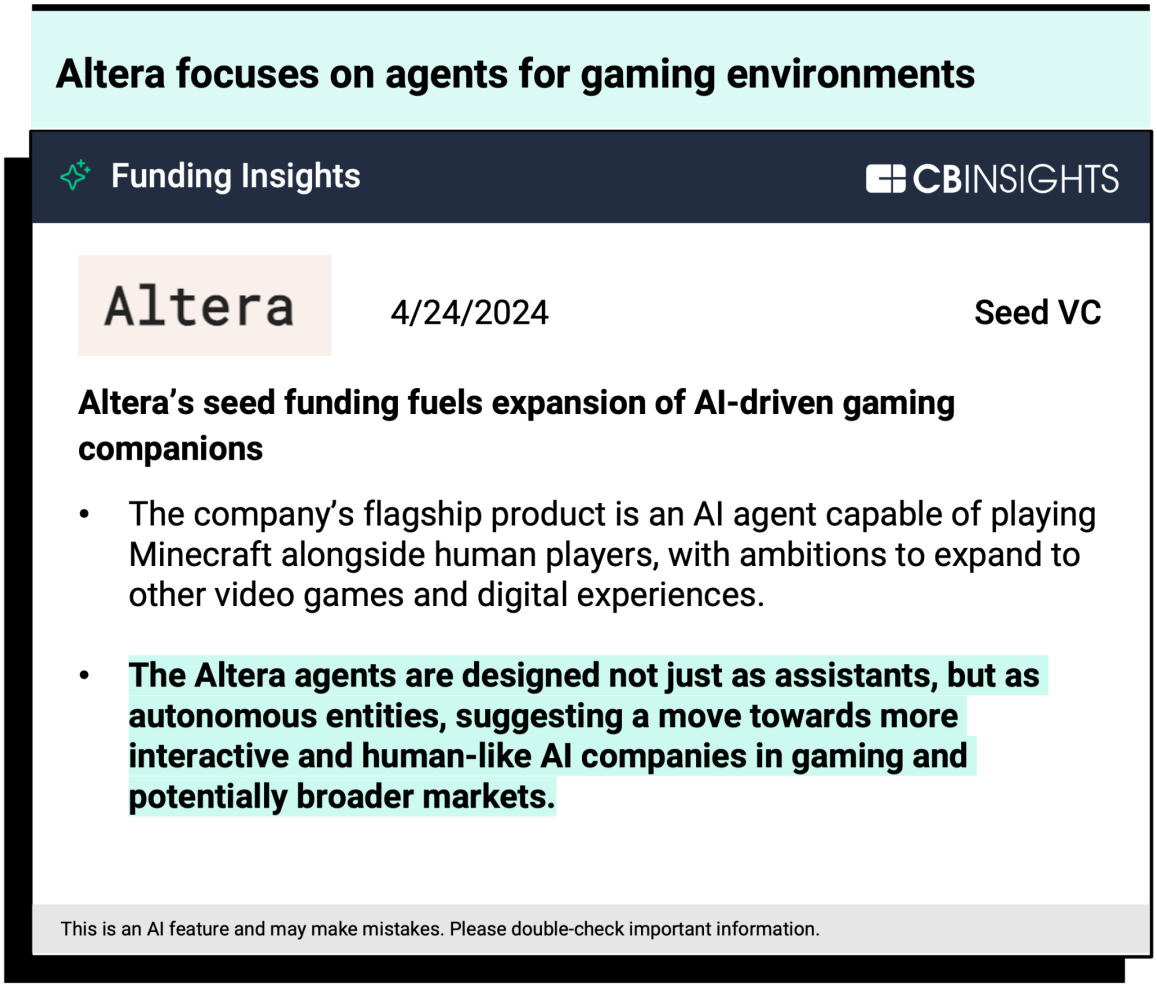

Source: CB Insights Altera Funding Insights

Expect to see far more activity and research related to gaming, given it offers a pre-made digital playground already optimized for humans to engage with.

Developments here may have implications for the broader AI agent space. As DeepMind highlighted in March, "Video games are a key proving ground for artificial intelligence (AI) systems. Like the real world, games are rich learning environments with responsive, real-time settings and ever-changing goals."

Looking ahead

More companies currently in the workflow automation and copilot/assistant space will move toward becoming (or at least trying to become) fully autonomous solutions.

At the same time, given that many of these agents are unproven — and still largely unreliable — we remain years away from widespread enterprise adoption.

Lack of trust will be the biggest barrier to adoption. In response, companies like Vijil — which raised $6M in funding from investors including Google's Gradient Ventures in July 2024 — are building security solutions to evaluate AI agents' behavior and trustworthiness.

Eventually, LLMs may become sophisticated enough to solve the reasoning and planning hurdles that agents currently face. Though another intelligence architecture altogether may be necessary to solve the "non-linguistic aspects to thinking and reasoning," as Sequoia Capital's Sonya Huang and Pat Grady write, to achieve higher accuracy.

In the meantime, expect some shaking out of the space — the most notable recent example being Amazon's acqui-hire of Adept's leadership. Other companies like Imbue, which raised over $200M in 2023, have yet to release a product.

But the promise of an empowered digital workforce remains — a trend that would reshape industries as we know them. As AI agents proliferate across enterprises, multi-agent architectures will emerge to manage them effectively, such as having a lead agent coordinate with specialized sub-agents. The implications would be enormous, changing how companies hire and scale, not to mention how they interact with each other.

Enterprises must pay attention to these developments — or face the risk of being outcompeted by companies that do.