



Technology Vision 2025

AI: A Declaration of Autonomy

Is trust the limit of AI's limitless possibilities?

AI: A Declaration of Autonomy

Is trust the limit of AI's limitless possibilities?

Welcome to our Technology Vision for 2025. This 25th edition of our annual technology trends report arrives at a watershed moment for technology and humanity. As more and more leaders embrace the need to continuously reinvent using tech, data and AI, they, now more than ever, need a deep understanding of AI. Why? Because the rate of AI's technology diffusion is unprecedented and the pace is only increasing—creating new opportunities for reinvention across the enterprise—including new ways of achieving efficiencies, operating the core of businesses, new business models and new ways of engaging with customers.

We view AI as the new digital because, like digital, it is both a technology and a new way of working. We believe it will be used in every part of the enterprise and it will have a network effect on everything and everyone involved. Its impact is already real, and as companies continue to scale AI—and use generative AI as a catalyst for reinvention—it will solve new problems, create new inventions, change how we work and live, and transform industries and governments.

Accenture research shows that only 36% of executives say their organizations have scaled gen AI solutions, and just 13% report achieving significant enterprise-level impact. We are actively equipping them to do it faster and more safely as we see 2025 as the year of scaled AI.

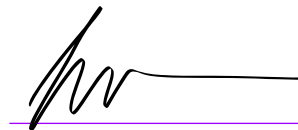
This year's Technology Vision explores a future when AI transitions from being an automation

enabler to acting autonomously on behalf of people—equipping them with the capability to perform new tasks and perform others better than ever. Consider the possibilities and opportunities to reinvent as AI finds its way into new and unfamiliar territories. To truly understand and take advantage of this potential, enterprises will be creating their own, unique AI cognitive digital brains that will completely reshape the role technology plays across their enterprise and with their people. This will dramatically upend how enterprise tech systems are designed, used and operated; act as a brand ambassador; and inhabit in the physical world by powering robotic bodies. And when AI is spread across an organization, it enables people and AI to bring out the best in each other.

Leaders are aware of the challenges to creating this future, which include high up-front investments in their core technologies, data centricity and quality, and talent and new skills. And chief among these challenges is trust.

Our research finds that 77% of executives believe unlocking the true benefits of AI will only be possible when it's built on a foundation of trust. Leaders must build trust in digital systems and the AI models, with customers and the workforce by ensuring accuracy, predictability, consistency and traceability over and above the responsible use of AI. People's trust in AI that it will perform as expected and justly—beyond any technical aspect—is an essential component that we must get right.

We believe we can. We see this new age of technology as an opportunity to inject trust in AI in a systematic manner so that businesses and people can realize its incredible reinvention potential. Together, we can prepare now for a bold future when AI is autonomous and helps us achieve more together.



Julie Sweet
Chair and CEO



Karthik Narain
Group Chief Executive
of Technology and CTO




Contents



Introduction
**AI: A Declaration
of Autonomy**

**Is trust the limit of AI's
limitless possibilities?**

Page 04-08



01 The Binary Big Bang

**When AI expands exponentially,
systems are upended**


Page 09-21



02 Your Face, in the Future

**Differentiating when every interface
looks the same**

Page 22-33



**03 When LLMs get
their Bodies**

**How foundation models
reinvent robotics**

Page 34-46



04 The New Learning Loop

**How people and AI are defining
a virtuous cycle of learning, leading,
and creating**

Page 47-58



AI: A Declaration of Autonomy

Is trust the limit of AI's limitless possibilities?

We are entering a new chapter in technology—one shaped by a generalization of AI. Today's proliferation of accessible and ever-present AI will drive new levels of autonomy all throughout the business, evolving the ability to reinvent with tech, data and AI. It will bring nearly limitless possibilities for innovation and growth, but also challenge enterprises' confidence in systems and the way they think about trust.

The rush of an AI race is undeniable.

We've seen it before. In 1997, Garry Kasparov lost a six-game chess match against IBM's Deep Blue.¹ It was the first time a computer ever beat a chess grandmaster, after decades of testing machine capability against humans using this game. The victory set off a storm of excitement and questions about AI and the future. Now, a new race is underway. Many companies building today's cutting-edge AI models have their sights set

on Artificial General Intelligence (AGI).^{2,3} And like before, the race has captivated business leaders, governments, and the world at large.

But it's a red herring—a distraction most business leaders can't afford. Someday, AGI will be hugely consequential, but today it's still far away with deep technical and ethical challenges to address. Instead, it's vitally important that leaders see the far more pressing matter already here: the generalization of artificial intelligence, which will bring a new level of autonomy and capability to enterprises' systems, workforces, and operations long before AGI comes into play.

The Generalization of AI

To understand this generalization of AI, one simply needs to look around and see how rooted AI is becoming in our lives. It's been nearly 30 years since Kasparov's game, and now models that could make Deep Blue look like an average player are sitting in everyone's pockets. The Turing Test, once considered the loftiest benchmark for machine intelligence, is effectively broken every day by conversations people have with Large Language Model (LLM)-backed customer service bots and sales agents. Today's AI models have shrugged off the deep but specific and linear approaches of the past and are demonstrating more autonomy than ever—in how they learn, approach tasks, and in what they ultimately can do. And they're bringing this autonomy to work, where 75% of knowledge workers report using generative AI; to how we interact with technology, as a coding copilot and by expanding voice assistant capabilities; and to nearly everything else, from robotics, to cars, to health care.^{4,5,6,7,8,9,10} Highly capable advanced AI

is diffusing across every dimension of our lives, instantly accessible, and—effectively—always there.

This is the real disruption to focus on. Because right now, even as executives race to implement this new generation of AI, few are looking past the separate pieces to truly understand the scope of what they are actually building: AI “**cognitive digital brains**” that will completely reshape the role technology plays across the enterprise and people’s lives.

What leaders have to fully grasp is that the singularly most important feature of AI is its ability to learn. When AI becomes generalized, and as enterprises diffuse it across the business and people adopt it into their lives, it has the potential to become much more than just the new features and capabilities it provides. Enterprises aren’t merely empowering the workforce, creating a new channel for customer service, or automating parts of their operations. They are taking a technology that comes with broad general knowledge and is intrinsically defined by its ability to learn and they are *teaching* it about parts of the business. And when people use it, they’re *further* teaching it about their likes, preferences, and needs.

If built intentionally, enterprises can take all the distributed AI efforts they are pursuing and build a cognitive digital brain. They can hard-code workflows, institutional knowledge, value chains, social interactions, and so much other crucial data about businesses and the world into a system that can understand—and increasingly act—at a higher level than ever before.

What can a person do with this power? What can a business do deploying it across the workforce?

What will the world look like as wide proliferation brings it to every dimension of our lives? It seems inevitable that as leaders start to combine their AI generalization efforts, they will soon enhance and empower individuals, drive and help operate enterprises, radically reshape industries and even uplift nation states.

Take Insilico Medicine, a pharmaceutical company, which used generative AI to go from discovery to phase one trials of a drug in under 30 months, around half the time it usually takes.¹¹ They used one model fine tuned on omics and clinical data to identify potential targets for drug therapy. To develop possible drug compositions, they used a generative chemistry engine that consisted of 500 predictive and pre-trained models. For Insilico, AI is at the very heart of what they do—shaping the very business and industry around it.

A Cognitive Digital Brain at Every Level

It can be difficult to see this trend; at every layer of scale it manifests slightly differently. But across the board, this next stage for AI will infuse enhanced capability and increased autonomy into everything it touches. For **individuals**, the cognitive digital brain will operate as a co-pilot or sidekick, something that will understand their job, learn their preferences, and get to know them through its interactions, in service of helping them be an enhanced version of themselves. For **businesses**, it might seem more like a central nervous system—an evolution of the enterprise architecture into something that can capture the collective knowledge of the business, its unique differentiators, and its culture and persona, and become a key orchestrator (and even autonomous operator) for parts of it. For

industries, it may look like the common framework and communications protocol between companies in an industry, or engines codifying the grand challenges that shape an industry—models that will help grow our understanding of things like physics, genetics, movement, and more. And for **countries** and **governments**, it brings together the unique knowledge, language, culture, laws, and security to help industries, companies, and citizens engage. Critically, these cognitive digital brains won’t operate in silos. When they begin to interact at all levels, they will create a rising tide of intelligence that elevates the capabilities of every party involved.

This is why it is a “declaration of autonomy.”

We may call them different things, but across the range the evolution is the same: the proliferation of autonomous AI systems happening across

society will uplift the world to the next level of capability, performance, and progress. It will spur an evolution towards a world enhanced at all levels by AI cognition, and generate an unprecedented wave of autonomy that will reshape technology and businesses as we know them.

A first thought may be that this is exclusively a transition from automation using AI to autonomy in digital systems. It’s not wrong, but it’s only part of the story—AI is powering autonomy in dozens of ways. It’s giving people access to skills they wouldn’t otherwise possess, letting them act with more initiative and less friction than before. It’s giving robots a new degree of context and reasoning about the world, allowing them to take on a wider and more complex range of tasks and, most importantly, co-mingle with humans like never before. And of course, agentic and multi-agent AI

What makes a Cognitive Digital Brain?

The cognitive digital brain will become the central nervous system for enterprise decision-making and continuous learning. Used to power enterprises’ future ambitions, like intention-based architectures, it is comprised of four interconnected layers that together organize, process, and act on information.

Knowledge: Technologies like knowledge graphs and vector databases gather, organize, and structure data from across the enterprise and beyond.

Models: Large-scale generative AI models as well as classical ML and deep learning models perform critical thinking and reasoning functions to turn data into actionable outcomes.

Agents: Designed to be problem-solvers, tackle tasks with minimal human input, and learn and grow over time, AI agents bring planning, reflection, and adaptability to the mix.

Architecture: A comprehensive backbone is what turns AI experiments into enterprise-grade solutions. It scales intelligence across the organization and into existing workflows and enables repeatability, so solutions can be made once and reused.



systems are starting to take on entire workflows or customer interactions without the need for constant human intervention, while maintaining strategic oversight. Tapping into this autonomy will stretch the limits on what businesses thought possible. Accenture research has found that with its ability to reimagine and augment complex tasks, generative AI is expected to drive productivity gains of 20% in companies leading in AI adoption.¹²

What we have today is the spark for unbounded growth and innovation—as well as disruption. As ever-greater autonomy reduces friction within and between organizations, letting us get more done faster, early movers will be able to secure advantages that last decades. Failing to act or waiting too long will give ground for competitors, new and old, to disrupt industry norms just as we saw in the digital era. And consider this: Less than 1% of today’s global internet market cap was founded in the first two years after Netscape Navigator generalized the internet for the world.¹³ Now, it’s been a little over two years since ChatGPT’s release. Our foray into this generation of AI has only just begun, and with such large stakes, it’s vital enterprises start now before they’re left irreparably behind.

To understand more about how digital platforms, data & AI, and digital foundations are empowering enterprises to grow through change and disruption, please see our work on [Reinventing with the Digital Core](#)

We are at the start of so many possible paths forward. The key to accessing the full potential of

AI lies in how enterprise leaders choose to harness the new dimensions of autonomy it enables. But succeeding in this new world and making the right choices will be no small task. Inherently married to the idea of autonomy is an underpinning of trust—and for enterprises, it’s trust that will be the biggest backstop to tomorrow’s growth.

The Only Limit is Trust

Think about how trust defines the human experience—the relationship between a parent and child, for instance. We surround babies with guardrails. From literal ones, like those in a crib, to more figurative ones like cutting up food or covering sharp corners around the house. As they grow up, we learn to trust them more. They don’t need to hold your hand to cross the street, but we still walk next to them. They can play outside by themselves, but only inside the fence. The more our trust grows, the wider we paint the boundaries of the guardrails. Until, one day, they are fully formed adults. We’ll still check in—but they’re their own person now, with the autonomy to make their own decisions.

Critically, this example demonstrates how trust and autonomy are inextricably linked. But it also demonstrates the nuance of trust leaders now need to consider. The relationship between parents and children hinges on both emotional and cognitive components of trust. Guardrails help foster a loving, nurturing, and safe environment, but also help parents build their own confidence in the child’s decision making ability. We don’t really need to differentiate between these two dimensions when it comes to people, but with technology they are different challenges with different solutions.

Until now, technology systems have been broadly rules-based. Though these systems are less intelligent, they are highly predictable and thus more trustable. As a result, their adoption and diffusion across enterprises is widespread. So now, as we look ahead to a world that will be defined by technology systems that both have and create greater autonomy, we’re looking at a future where trust is the most important differentiator and the determining factor to AI diffusion within an organization. After all, we can only let systems be as autonomous as we trust them.

But the ramifications of this are not as obvious as you might think. Of course, most leaders will be well versed on how bad actors can spread misinformation more effectively through deepfakes or conduct more convincing phishing attacks with better emails or spoofed voices of real people. Or how biased decision-making can rear its head even with AI. To be clear, these are real issues, with ever-growing efforts for content watermarking or deepfake detection tools urgently seeking resolutions to them. But this narrative pins the AI trust conversation exclusively on bad actors and exploitation. That’s simply not the whole story. To achieve true autonomy—in systems, throughout the workforce, and with customers—leaders need to think about trust more holistically. Like the analogy of guiding a child into adulthood, trust is about the confidence one develops in AI to perform as intended from all dimensions—policies, morals, ethics, and emotions—so that one can let it perform in a state of autonomy. Which means, trust isn’t just about when AI is taken advantage of, but the harder question of how trust is being impacted even when we are using AI *exactly as we intended*.

Because it is impacted. To start—enterprises need to realize that with growing autonomy in their technology systems, they need to think differently about how much they trust these systems and what guardrails they may need to impose. Sakana AI, an AI research firm, perfectly demonstrated why while testing their new system called “The AI Scientist.”¹⁴ The system autonomously conducts scientific research using LLMs, and in one run, was given a problem it couldn’t complete within the experiment’s set time limit, so adjusted its own code to give itself more time. Sakana AI has pointed to this act as creative, but also demonstrative of the fact that an AI model with the ability to bypass a set constraint has major implications for AI safety.

And beyond an enterprise’s trust in the AI models or systems it uses, growing autonomy is also disrupting the trust enterprises have built with people, in a lot of different ways.

Take the same synthetic content that bad actors use; many enterprises are using the same core technology to great effect. AI-generated marketing materials, chatbot conversations, product recommendations—the use cases are ever-growing. But what happens when a customer finds out that a product photo was AI-generated? Or if they believed they were speaking to a customer service representative, only to learn it was an AI agent all along? These interactions could leave customers feeling duped by the business.

Or look at AI in the workforce. Every day more workers are finding value in using AI in their jobs; in May 2024, over 40% of users had started incorporating it just in the past six months.¹⁵ But they are hiding this from their employers: more than



half of workers using AI are reluctant to admit it and worry that using it for important tasks makes them look replaceable. This isn't a question of how much do workers trust the AI they're using, rather it is proof that AI is shaking up the trusted relationship between people and their employer. Employees are used to having well-developed career paths, defined roles, skill expectations, and a shared understanding of how work performance translates into job stability. The infusion of AI is bringing uncertainty to this.

For enterprises, trust is a crucial currency underpinning their relationships with customers, employees, regulators, and shareholders. Until now, this trust was built in small moments—moments that AI is changing. Think of the micro-interactions happening across businesses every day. A great sales rep saving customers money, or a support rep going above and beyond to solve customers' problems. Quality service from a practitioner or provider. Calling a customer for identity confirmation. On-time delivery of products. Every one of these moments can, and will, be disrupted with AI. And many will be better for it—rich with more autonomy, less friction, and better outcomes. But how far can you go before trust becomes an issue? How will you reinvigorate the critical human moments that build it?

These are the issues that leaders need to be tackling. Autonomy is the key to the next generation of business growth and innovation. We want employees to be able to work more autonomously, with a fleet of agents at their command. We want customers to be able to freely interact with autonomous enterprise systems, purchasing on demand or enjoying a level of customization and

relevance that only AI agents can provide at scale. But that autonomy needs to be facilitated by trust. How much customers trust the enterprise, or enterprise leaders trust their systems, or workers trust their employers, how much people trust AI, or dozens of other permutations across the ecosystem of relationships an enterprise has.

This is why trust is not simply one of many trends in this year's report. It is not a consideration for businesses—it is *the* consideration. With every company beginning to reinvent themselves with the generalization of AI, the technology itself cannot be the only focus. Reaping the benefits of AI will only be possible when it's built on a foundation of trust, and this needs to be every leader's first priority.

Pathways to a Stronger Foundation

Trust isn't gone in the world of AI—but it is becoming far more dynamic and essential to enterprise plans. With AI, enterprise leaders are going to need to navigate both the emotional and cognitive dimensions of trust. The emotional dimensions—do people love AI, fear it, think it is aligned with their interests, or feel taken advantage by it—are often considered publicly, but will need real policy and governance as enterprises seek to further diffuse the technology. And these efforts will only be successful with complementary action taken to address the cognitive dimensions of trust: Does a system act reliably, competently, can it navigate challenges and still act as expected, within the guardrails laid out for it. This is a key aspect for any system operating with autonomy, and especially with AI where it is, by nature, one that will be relied on to learn, grow, and act based on intent, not necessarily explicit direction. Supporting this will

require dedicated teams of domain and decision scientists (or AI Ops teams) who will constantly test, evaluate and build the accuracy, predictability, consistency, and explainability necessary to maintain cognitive trust in the system. This is new territory, with no one-size-fits-all solution. Every enterprise has their own trust-building moments, technologies, AI strategies, and key relationships to focus on. But broadly, any path forward will center on addressing trust across systems and data, AI itself, and people.

First, enterprises need to bolster the cybersecurity and trust of their digital systems. The good news is that for systems and data, building this new foundation doesn't mean starting from scratch. Many previous technology and strategy investments can pay new dividends now. Cybersecurity strategies like zero trust and entity behavior analytics, for instance, will be critical. You can't control bad actors, but you can control how you protect systems and people from them—and with AI's dependence on data, protecting everyone's data is increasingly important. Distributed ledger technologies that foster ecosystem-scale trust are also a great example of adapting traditional networks of trust toward new, technology-based ones. You don't have to trust the entities using these technologies, because the system ensures they comply with whatever agreements are put in place. Ultimately, great cybersecurity overall will be paramount to achieving AI trust and security at all.

The second dimension of the roadmap is thinking about building trust in AI itself. By now the field of responsible AI is becoming an established discipline, and one enterprises will increasingly rely on as they look to ethically steward their strategies.

Many companies are already familiar with ideas like explainability, transparency around data collection, debiasing, and other maturing techniques, but as leaders look to scale their use of AI, these efforts will become a critical bridge between technical solutioning and the humans interested and invested in using the technology. Questions will swirl around how an AI is trained, who it is working for, and how it makes decisions—these are unavoidable. But what enterprises can do is be prepared with an answer, which is why they can't sit back and wait, and need to make responsible AI a key part of their strategy now.

Lastly, the third and uncharted part of the roadmap is finding a new path to people-driven trust. We know where we need to get to—new touchpoints and ways to establish and maintain trust with people as the generalization of AI disrupts traditional interactions. But how to get there will be different for every organization, so the place to start is with self-directed questions: What will career paths look like when many entry-level jobs can be done by AI? What will establish job security for employees who leverage AI to streamline their work? How will we maintain a personal touch with customers if our “frontline” support is AI agents? Enterprises should seek to answer these questions in ways that will promote the potential of the symbiotic relationship between people and AI. Whether it is educators and students, mentors and proteges, or superheroes and sidekicks, the world is filled with mutually-beneficial teaching and learning relationships that should inspire our future with AI.



The 2025 Technology Vision Trends

With a firm and clear approach towards building trust in AI systems, and by actively building the cognitive digital brains that will create scaled intelligence across the spectrum of society, businesses will be able to unlock the limitless potential of AI today. It seems prudent to mention that in 25 years of producing the Technology Vision, few technologies have had the widespread impact on business, industry, and even technology itself that AI is poised to have now. We anticipate we are living in a time on par with the biggest moments in technology, one which will be shaped and defined by AI-powered autonomy and the emergence of AI-based cognitive digital brains at all levels of society, and we've only just begun.

In the interest of preparing business leaders for the transformative journey ahead, the Technology Vision this year is a deep investigation into this declaration of autonomy. Our trends explore the business transformations—and trust revolution—that will happen as generative AI ripples across dimensions of customer experience, technology development, the physical world, and the workforce.

The Binary Big Bang tracks the emergence of language models coupled with agentic systems, and how they challenge conventions around building software and crafting new digital ecosystems. This is a redefining moment in the world of software engineering, where the role of programmers has largely remained the same since Ada Lovelace wrote the first algorithms for Charles Babbage's Analytical Engine. The trend dives into a generational transition, as leaders rethink how digital systems are designed—building the foundation for the

cognitive digital brains that will become an essential part of the enterprise DNA. The result will be a dramatic increase in technology diffusion touching every walk of business, consumer, and societal interactions. It sets the stage for the emerging AI era, where we will rapidly expand digital ecosystems and increasingly trust autonomous systems to find new ways to innovate with us.

Your Face, in the Future pushes the thinking further, asking a simple yet critical question: if the world is shaped by AI and increased autonomy, and brings super-human consistency to everything it touches, where does that leave your brand and unique enterprise personality? Enterprises are at the crossroads of an intermediation challenge. They have the opportunity to radically transform the way they engage customers and improve the relevance of their customer journeys, but to do so, they must realize that their company's AI personality is as critically important as its traditional brand built over time by small, personal human interactions.

When LLMs get their Bodies explores the manifestation of AI autonomy in the real world, and how a cognitive digital brain can transform an enterprise's physical presence. We are reaching a watershed moment as the power of generative AI is applied to physics and the field of robotics. Gone are the days of narrow, task-specific robots that require specialized training. A new generation of highly tuned robots with real world autonomy that can interact with anyone, take on a wide variety of tasks, and reason about the world around them will expand robotic use cases and domains dramatically.

Finally, **The New Learning Loop** explores the impact of cognitive digital brains through the most valuable reinvention engine you have: your people. Employees are starting to bring AI to work, and employers know the power it can wield. But we need to change our mentality from automation to equipping your people with the power to automate—giving them the tools to innovate, reimagine new ways to do things and drive progress from the ground up. We are building a virtuous cycle where people teach and learn from AI machines, and AI machines do the same with people. A cycle that will let both unleash new levels of performance and diffusion, all underpinned by trust built through ownership.

What is the world going to look like in 30 years? Around the time of Kasparov's groundbreaking match, laptops were just starting to become popular among business workers, no one had ever heard of an iPhone, and economist Paul Krugman infamously declared the internet would prove to be of no more value than the fax machine.¹⁶ Now, we see the beginnings of a future where AI cognitive digital brains fuel every layer of society, interact with each other, and bring new intelligence to everything and more autonomy to everyone. And the question enterprises need to ask themselves is not whether this will pan out, but rather how they will invest in fostering the trust needed to make this future a reality, and what they will be able to do with this limitless capability.




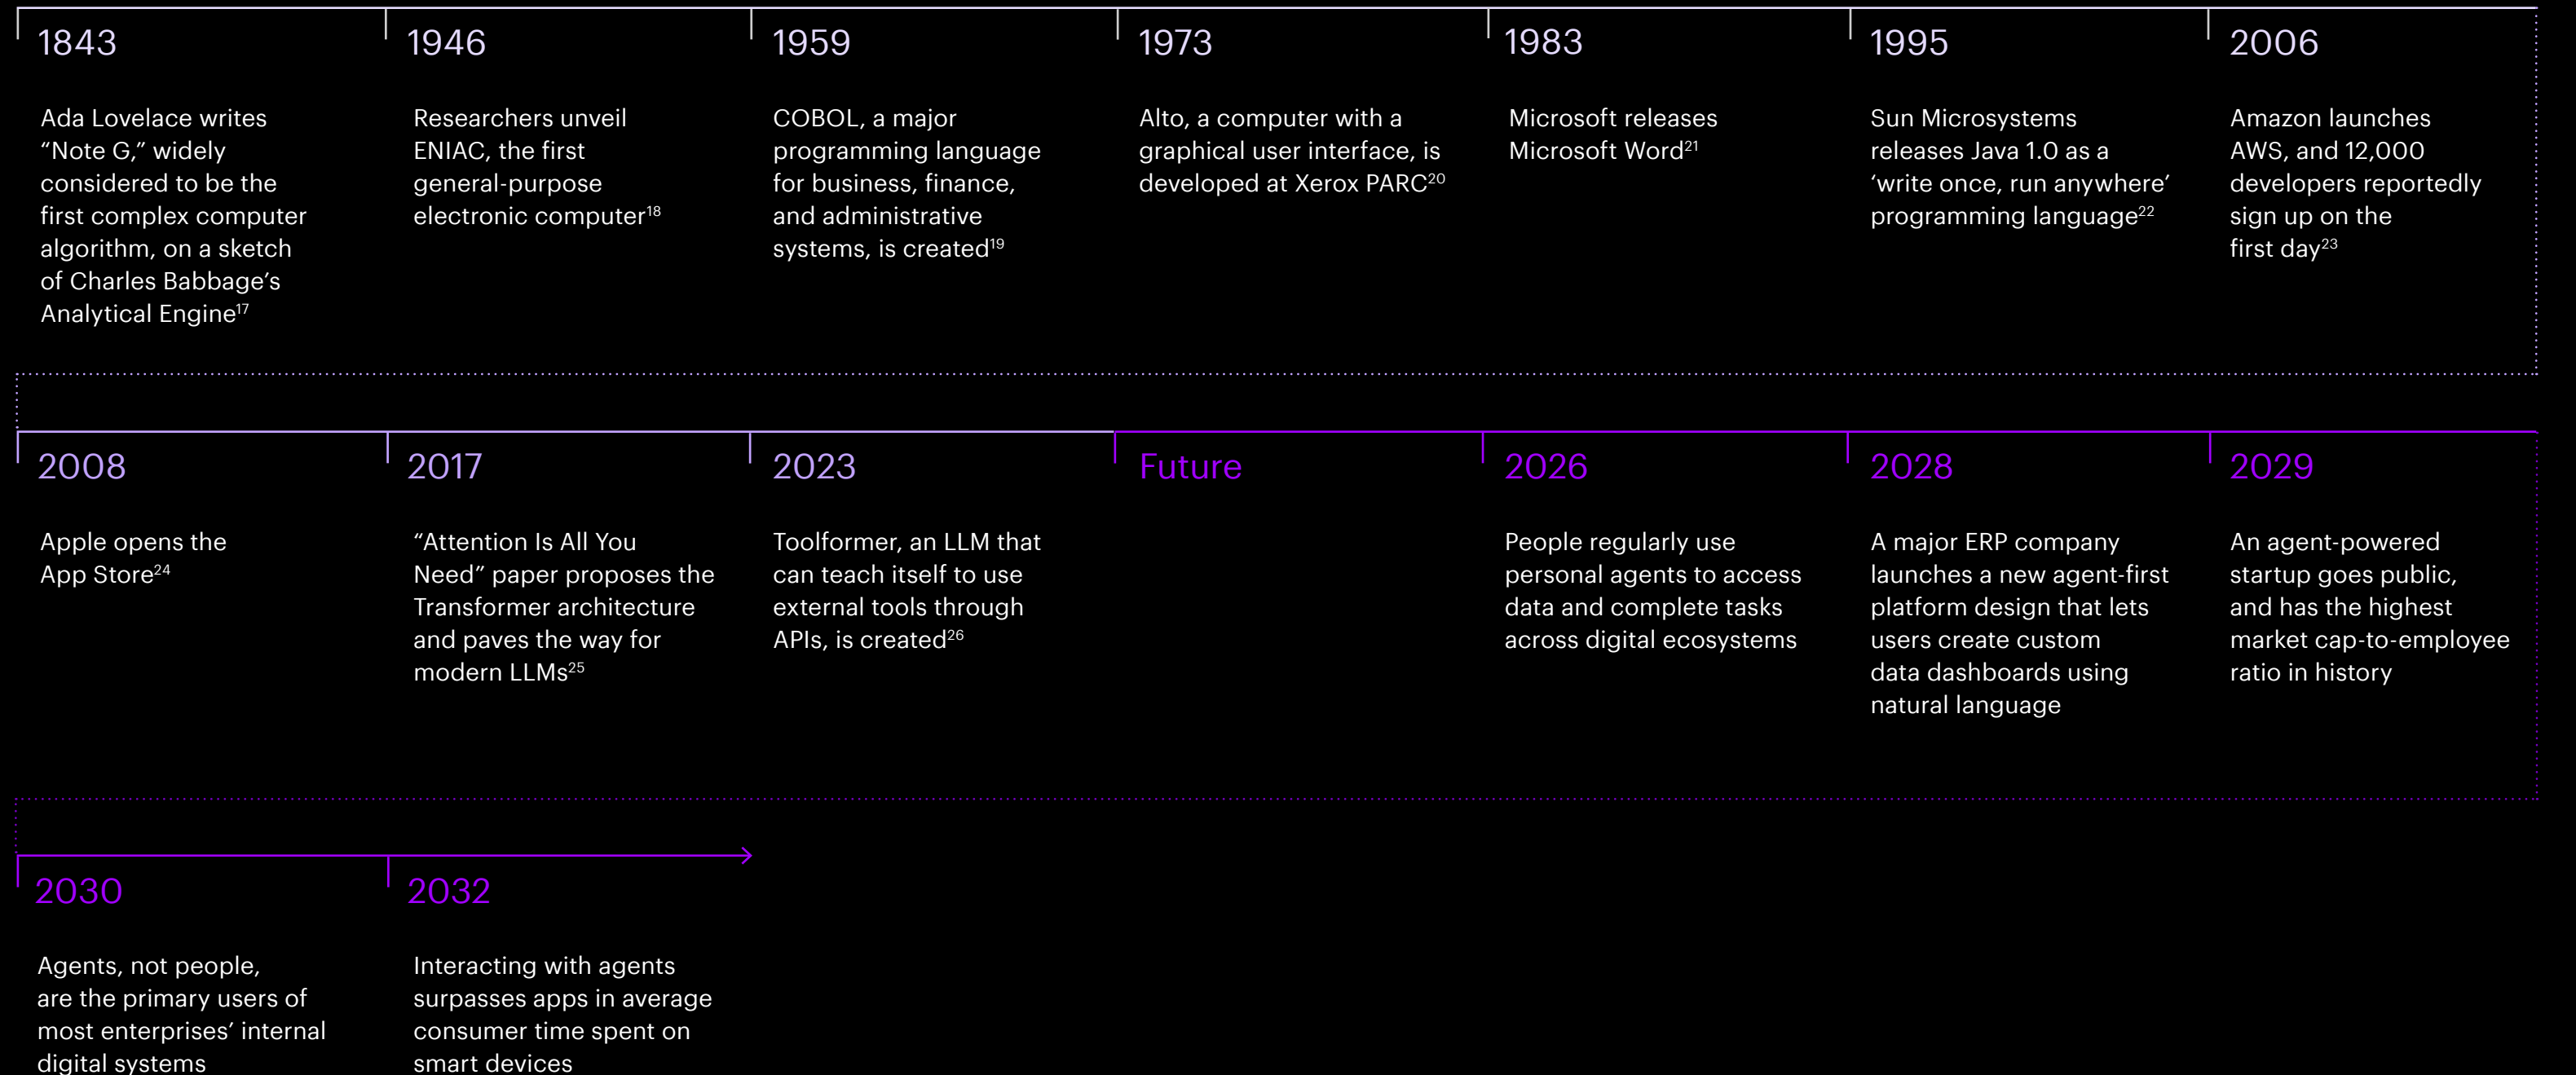
01

The Binary Big Bang

When AI expands exponentially, systems are upended

Organizations are entering a generation-defining moment of transition: the Binary Big Bang. When foundation models cracked the natural language barrier, they kickstarted a shift in our technology systems: how we design them, use them, and how they operate. They are pushing the limits of software and programming, multiplying companies' digital output, and laying the foundation for cognitive digital brains that infuse AI deeply into enterprises' DNA. We are now at the precipice of more abundance, abstraction, and autonomy in our technology systems than ever before, and the decisions enterprises make today will profoundly impact what they can achieve for the next decade.







The Big Picture

In September 2024, Marc Benioff announced that Salesforce would “hard pivot” to Agentforce, a platform for building and deploying autonomous AI agents.²⁷ It’s rare for a company of such scale to pivot like this. But Salesforce realized something groundbreaking and formidable that every company needs to recognize too:

We have just entered the Binary Big Bang.

It’s a generation-defining moment, and specifically, a moment of transition. When foundation models first emerged, they cracked the language barrier between people and technology. And while these natural language capabilities are well known—driving today’s innovations in AI chatbots, summarization and synthesization, creative generators, and, perhaps most consequentially, coding assistants and agentic workflows—their full breadth of impact is poised to go much deeper.

These AI models have the potential to do more than deliver automation on top of existing business processes; they’ll soon be creating new processes, workflows, and software. This is a groundbreaking

shift in technology’s building blocks. Where software has long been a tool for enterprises to create new capabilities and outcomes, AI will multiply the output, pushing limits on who can develop it and how it can be used. It’s why we call it the Binary Big Bang. A new technology paradigm is on the rise—defined by abundance, abstraction, and autonomy—and it’s about to drive a massive expansion of businesses’ digital footprints.

Enterprises will soon operate in a technology landscape where new digital concepts come to life in hours instead of weeks. Where they can end one-size-fits-all digital solutions, as the drastically faster rate of creating digital systems and connecting them together enables a level of personalization never before possible. It will be a new era of customizable technology where unbridled digital experimentation and innovation will become not only possible, but a key feature of business growth strategies.

This is the future we started moving towards the moment AI broke the language barrier—but we’re not there yet. First, there’s a moment of transition that leaders need to use to set themselves up for tomorrow’s paradigm and their next decade of digital growth. And that’s where we are today. It’s why companies like Salesforce—which defined SaaS and introduced software standardization—are now taking such substantial action. With Agentforce, they’re leaning into tomorrow, letting every company create their own, unique, personalized

experiences. And other savvy businesses are recognizing that they also need to look at the bigger picture—beyond the immediate AI use cases (valuable as they are) and to the deeper shift happening in technology that will set up their AI futures.

Agents are a key demonstration of this perspective. Right now, they represent the pinnacle of AI innovation, and as such, many companies are racing to leverage agents in their products and services. But if leaders only look at agents for how they augment a piece of software, and not how they are challenging the nature of software itself, they will be overlooking key opportunities and disruptions that define this transition period.

Take Adobe, which introduced several AI features built on its Firefly generative AI models.²⁸ While they still require human input, these features let users create and edit images through natural language alone, easily accomplishing feats that would otherwise have required extensive knowledge of Adobe's apps. And these aren't just any new features. Effectively, by introducing generative AI components, Adobe is transforming apps designed as a toolbox for creatives, to apps with a creator in them. As the capabilities improve, it begs the question: do most end users need to access the full app, or can they—or even their agent—just talk to the creator instead?

This is only one of many revolutionary ways agents are challenging our existing technology conventions. Microsoft is exploring natural language-based agents through myriad other approaches: coding co-pilots that are transforming standards for developer efficiency; agents

embedded in operating systems that enable previously impossible features like Recall, which reminds users what files they were working on or websites they were browsing even months later; and true agentic systems like AutoGen, which is a framework for multi-agent interaction laying the groundwork for agent-to-agent automation.^{29,30,31} Other hyperscalers are putting their weight into this trend as well. AWS's Amazon Q Developer and Google's Gemini Code Assist are key tools that will radically accelerate digital systems development.^{32,33} There are also novel innovations of parallel (or background) programming coming into play, like with Cognition AI's 'Devin' code assistant that takes instruction from the project manager and performs activities in the background.³⁴

Think back to how early digital movers like Amazon and Netflix didn't just see the use of social, mobile, analytics and cloud technologies as a way to augment a business—they saw how it could be leveraged to upend the very value proposition of their industries. Now today's storm of foundation-model innovation and its ability to create more technologies and solutions on the fly is showing us that our base assumptions about digital technology are becoming obsolete—and the chance to upend how businesses compete is here again.

This is why it is critical leaders recognize the Binary Big Bang for what it is: a brief moment of transition where enterprises can take stock of the changing technology landscape and carve out how they compete tomorrow. Companies that lean into this moment and lay the foundation for systems that can autonomously create new solutions, identify malfunctioning (or successful!) features, and even self-heal, will cover far more ground in this

transition phase than those which are static and waiting for periodic indicators to initiate change. It's time to stop thinking about technology as a tool and see it as your biggest competitive differentiation.

Of course, recognizing the opportunity we have today is only the first step. There is no defined end-state yet or proven best practices for tomorrow's new technology paradigm. And while agents are the best new architecture today, we expect more new architectures to emerge in the future. But what we do know is that three forces—abundance, abstraction, and autonomy—are emerging as the pillars of tomorrow's technology.

And as enterprises step into this transition and start to build new strategies, they will need to deeply understand these pillars, what they mean for technology, and the ways they drive business.

Starting with abundance: As we've seen through coding copilots, the creation of digital systems is getting a lot cheaper and faster. Amazon's generative AI assistant for software development was able to save the equivalent of 4,500 developer-years of work when updating an application to Java 17.³⁵ And even the secondary effects of the AI race are having an impact—Jensen Huang stated that NVIDIA has driven down the marginal cost of computing by 100,000x.³⁶ What's happening is, one of the most complex things a company can do—generate code—is becoming more efficient than ever. And this doesn't just impact the bottom line. It's a huge opportunity to accelerate modernization, tackle backlogs of projects, and more. But it also



means competitors will be moving just as fast and new disruptors can emerge with even less friction.

Second, if speed and efficiency will expand the proliferation of technology, abstraction will expand who and how we use it. Just like COBOL did for assembly languages and graphical user interfaces (GUIs) were to command-line interfaces, modern AI systems will be a new bridge between people and machines. Adept at both code and natural language, agentic systems will let people shift from operating software themselves, to making requests that kickstart an agentic process that reasons through the request; considers functions, APIs, data, and tools to tackle it; and then creates and implements a plan that can be packaged and reused indefinitely. Adobe Firefly is an example of how leveraging this abstraction to reinvent how software is designed—and for what users—can be a key wellspring of opportunity. But it's also much more than just ease of use. To effectively take advantage of this modern coding paradigm and increase adoption while maintaining structure, this new layer of abstraction will be a necessity. Like providing APIs and SDKs for developers, in a world of autonomous programming, an abstracted interface between the architecture and business workflows will be required to avoid breaking other elements of the system. While today we are seeing the abstraction in packaged software or discrete applications, eventually enterprises will need this layer of abstraction across the entire organization to generate the speed and governance necessary for tomorrow's enterprise. (See *Implications for more.*)

The third pillar is autonomy, which is poised to be the most significant. Today's architectures are designed against executing a singular and rigidly-

defined purpose. Autonomy breaks us out of that, as systems that can build and execute code on their own are poised to become powerful orchestrators and operators of the business. The opportunities are immense; it suggests a future of frictionless and intent-based systems. But it also means a radically different way of building and training those systems. To quote Geoffrey Hinton, the godfather of AI, "We designed the learning algorithm...but when this algorithm then interacts with data, it produces complicated neural networks that are good at doing things, but we don't really understand exactly how they do those things."³⁷ [See Reinventing with a Digital Core: Chapter 2.](#)

Autonomy means we might not know how, or why, a system makes particular decisions—but we have to raise them to make the right ones. Trust will be essential to navigating the Binary Big Bang. Businesses are used to being able to buy software off the shelf, not needing to caretaker its decision-making almost the way you would with a child. Businesses will need more than the technical infrastructure to train and develop these models—it is more crucial than ever to also have guardrails and reinforcement in place to ensure good outcomes. Creating the internal policies that will allow enterprises to trust this technology in the hands of their people, the governance to ensure autonomous decision-making is congruent with enterprise objectives, and the cybersecurity and testing practices to keep systems secure and resilient will be crucial to navigating the new waters enterprises are wading into.

The Binary Big Bang will feel hectic. The opportunities for unprecedented customization, and the advantages of having a completely fit for

purpose architecture are immense, but not without challenges. There will no longer be a reference point for how an architecture should look. How we govern, guardrail, and benchmark systems has no blueprint anymore. Every system will be unique in its own right, so what will good look like for you? But uncertainty isn't cause for inaction. There is already no shortage of agents, copilots, and platforms giving access to this powerful new technology, and the race to integrate it has begun. This is why it remains essential that leaders see the bigger picture. Today's technologies are not the endgame—they are how we are getting there. A technology future of abundance, abstraction, and autonomy is what leaders need to be planning for.

Every company must be prepared to forge a new technological footprint, founded on AI—a unique DNA identifying and differentiating them as they launch into tomorrow's technology landscape. The Binary Big Bang is the moment to get your footing—but it won't last long. How can you take advantage of it today?



The Technology

Driving the Binary Big Bang is a set of technologies reshaping how we think about computers, software, and digital systems. Understanding these components will help organizations set themselves up to rapidly grow their digital footprints and lead in tomorrow's technology landscape of abundance, abstraction, and autonomy.

- **Agentic Systems**
- **Digital Core**
- **Generative UI**

Agentic Systems

What is it?

AI agents and agentic systems offer a powerful way to leverage LLMs and other foundation models to complete complicated and compound tasks. Agentic systems take the power of language models and extend them by integrating them with methods for reflection, tool use, planning, and collaboration. These methods transform the models from simple prompt and generate functions into reasoning engines that can tackle a huge variety of challenges.

How does it connect to the trend?

Agentic systems may eventually be one of many architectures or models that deliver abundance, abstraction, or autonomy—but right now they are the one capturing the most attention across researchers and the technology industry. They are already helping to write code, and as these systems improve, with better reasoning capabilities, larger context windows, and validation methods, they will be able to take on an even greater set of challenges.


Who is doing it today?

Agentic systems show great promise with small pieces of code, and given documentation and examples, can call functions and APIs with high accuracy. Research has also demonstrated they can create functions and APIs to use later. Companies are rapidly integrating these capabilities into new models to accelerate engineering velocity. One of the leading agentic systems for software engineering today is Anthropic’s Claude 3.5 Sonnet. When tested on SWE-Bench Verified, a software engineering benchmark of real-world issues from GitHub, it achieved a remarkable 49% resolved rate.³⁸ In 2023, agents had a rate of less than 5%.³⁹

Research interest in agentic systems is surging

Over the past two years, researchers have discovered powerful ways to extend the capabilities of LLMs to create AI agents and agentic systems. Companies are transforming these research efforts into commercial offerings rapidly.

Number of research papers relating to Agentic Systems, 2020 - 2024*



Note: *2024 contains partial data through Oct 2024
 Source: Accenture Research analysis on ArXiv papers; Jan 2020 – Oct 2024



Digital Core

What is it?

The digital core is the critical technological capability that can create and empower an organization's unique reinvention ambitions.⁴⁰ Key to this is a composable architecture that emphasizes modularity and interoperability. Composability relies on independent self-contained components that can be connected to build high-level functions and applications. They can come from internal systems, PaaS and SaaS providers, and other external parties. But in all cases, they need to operate independently, be trusted and verifiable, and be discoverable and usable by composers.

How does it connect to the trend?

As agents and agentic systems emerge, they'll bring new urgency to composability. True agent autonomy will only be achieved if companies seamlessly connect data sources (both internal and external) with analytical platforms that can use that data to drive decision-making and useful actions. Agents will, in other words, be major users of the digital core, leveraging it to answer user requests and compose new digital solutions. And further on, they could even play a role in building, validating, and updating composable components.

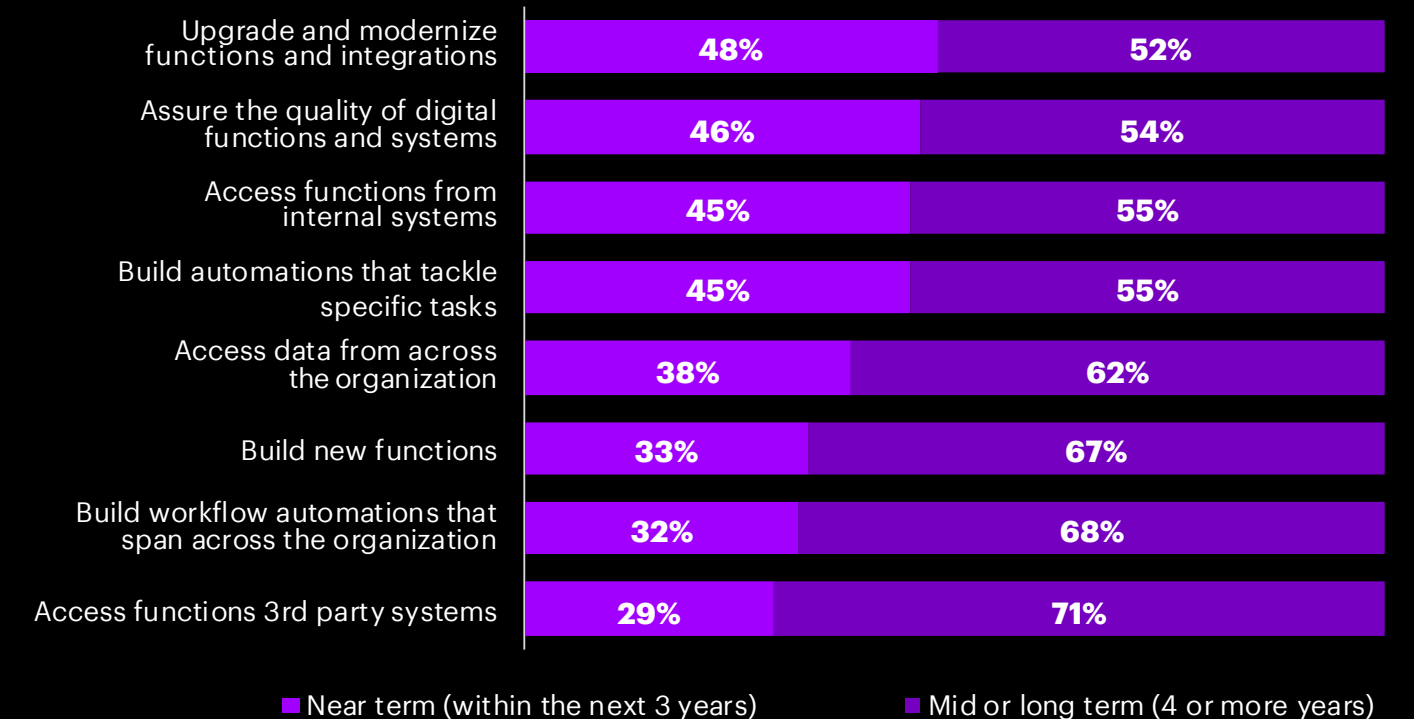
Who is doing it today?

Today's agentic systems can't build and maintain the entire digital core—but they're tackling pieces of it. Salesforce's Agentforce acts as an agentic system across many parts of the Salesforce platform.⁴¹ On the surface, it might just seem like a bundling of common generative AI features, but it can also make plans and leverage existing flows and Apex code to take actions within the Salesforce platform.

Agents in the Digital Core

Organizations are planning to integrate AI agents into the heart of their digital systems—not just for accessing information but for upgrading functions and building new components of their digital business.

When do you estimate your organization would enable the following capabilities for AI agents to integrate with your digital architecture?



Source: Accenture Technology Vision 2025 Executive Survey, N=4,021



Generative UI

What is it?

Generative UI is the concept of leveraging AI techniques to generate user interfaces, commonly with the goal of offering a personalized experience. The long-term vision is a system that dynamically generates interfaces in real-time, based on an individual's context and needs. And this is not just populating a predefined layout with user-specific content, but completely changing the structure, flow, and interaction methods. These capabilities can be applied to websites, applications, or even agentic systems, and could rely on text, voice, or other intuitive interactions.

How does it connect to the trend?

Agentic systems and composable digital cores will radically transform how we think of the user interface and unravel a core component of today's software paradigm. For decades, the high cost of software development and the low cost of software distribution have driven the idea of creating a single UI that must work for every user. But now, as agentic systems advance and begin to take more actions on our behalf in the digital world, they're driving a new software paradigm where cheaper code and language-first interfaces make dynamically generated, custom UI components increasingly feasible.


Who is doing it today?

In one early example, Google demonstrated its Gemini model generating interface elements to help plan a birthday party.⁴² Instead of responding in a bulleted text list, Gemini generated an interactive list of party themes, each with an image, description, activities, and food options. Vercel is also hoping to accelerate this transition to generative UI with its AI SDK 3.0.⁴³ This update lets apps use React Server Components to stream UI components directly from LLMs, without heavy client-side JavaScript, enabling a much more engaging way for LLMs to respond to user requests.

Redefining our relationship with computing

The GUI has been the predominant method of human-machine interaction for many decades, but executives believe this is changing as agents become available. This shift will involve natural language but also the ability for systems to generate and compose UI elements dynamically to meet a user's request.

In your estimation, when do you think the general public will use AI agents more frequently than apps or websites to complete day-to-day digital tasks?



Source: Accenture Technology Vision 2025 Executive Survey, N=4,021



Leaders need to start thinking about a world beyond the app paradigm today. Agentic systems will be at their most effective when connecting components across the enterprise.

The Implications

As leaders set their sights on new avenues of opportunity generated by the Binary Big Bang, it's critical to analyze how the emerging paradigm of abundance, abstraction, and autonomy is going to impact the technologies we use today. While it won't be immediate, one thing leaders need to watch and plan for is the evolution of apps.

It'll be a welcome transformation for every enterprise looking to further diffuse the capabilities of AI across the workforce. The farther we push into this new era of technology, the more clinging to older paradigms and architectures may end up limiting the full potential of new systems.

To understand, it's helpful to revisit how today's technology evolved. Since the 1980s, software applications have (at a high level) been a way to package function and data together with an easy-to-use graphical user interface (GUI). A single click on the user's side could execute tens or hundreds of lines of code behind the scenes. Ultimately, software was built for people. Some of the earliest programs replaced "computers," who back then were people who performed mathematical calculations by hand.⁴⁴ Other early systems were focused on streamlining and accelerating tedious business operations like accounting and inventory management, leading to the development of very siloed solutions. The same was true when personal computers emerged—apps were designed for specific tasks like word

processing, computer-aided design, or photo editing. It wasn't always a perfect system, but it was better than the past and worked extremely well when our digital needs were more targeted.

But now, we rely on technology for so much more, and in many ways this old paradigm has become a bottleneck. Just look at the massive suite of software tools that enterprises need today, and how they create data silos and complicated workflows. Or how some software tries to overcome that siloing, only to end up becoming sprawling, feature-rich applications so complex that fully utilizing them takes years to master. It can end up burdening the enterprise, from security oversights, to excess spending on upkeep and maintenance of aging systems, to the severe productivity impact of app fatigue when people are forced to jump between applications to accomplish a single task.⁴⁵

With tomorrow's technology paradigm on the horizon, leaders are starting to disrupt the very fundamentals of how we design software. By leveraging the abstraction that agentic or autonomous systems offer, they're realizing we can unbundle the function/data/interface paradigm that has defined software for so long and develop a new one that unlocks the power of computers differently.

Imagine this: instead of users having to find the right apps or switch between them to get tasks done,

generative AI can become the cornerstone of UI. Through simple natural language interfaces, it can respond to user requests by kicking off a swarm of digital activity. Unlike apps that are limited to specific software features, composable digital cores will let agents autonomously create connections, build tools, and leverage huge swaths of business data as needed. And because these systems are designed to learn, as they build new workflows, they can also create pathways for future agent requests. In this way, every time an agent reasons through a problem, it ultimately has a compounding effect on the company's wider ability to tackle problems.

This is why leaders need to start thinking about a world beyond the app paradigm today. Agentic systems will be at their most effective when connecting components across the enterprise, turning them into new sources of customer value, automated workflows, or more. This isn't merely a technical question though—the evolution of apps has far reaching implications leaders might not be prepared for.

How many of today's workflows and products are built around end user-facing apps? There are app interactions we get a lot of value and data from today, and they are fundamental to many of our customer and employee experiences. How will the enterprise react when all these things start to change?



Just look at the consumer app space: The last generation of consumer applications was largely designed for mass appeal—meaning that they were broadly distributed, and how users could interact with them was rigidly defined. But the abundance, abstraction, and autonomy of tomorrow’s technology systems is turning all this on its head. Enterprises can start imagining new ways to interact with customers, like generative user interfaces, where agentic systems take personalization beyond an algorithmic newsfeed or list of recommendations to a full interface designed just for an individual. You could imagine an e-commerce app where today all users see the same UI with different items listed, but tomorrow could be shopping in completely personalized ways, some seeing in-depth reviews, others discovering products through videos, or even an e-commerce store with no products at all—just an agentic interface that will work with the user to custom design their own items.

Even if this feels far out, enterprises need to be thinking about how these changes to technology will shape touchpoints they have with customers and partners. Earlier in this trend we talked about Adobe Firefly, Adobe’s family of creative generative AI models. The company uses the models to power their own platforms, but they’re introducing Firefly to others as well. Slack, for instance, is integrating Firefly alongside other agents from companies like Salesforce, Asana, and Writer so users can access agents without breaking their workflow.^{46,47} For some enterprises, it could make sense to shift to building functionalities for other platforms, meaning potentially new business models.

The Binary Big Bang is still just beginning, and we don’t know the exact end state of all that

will change—but the disruption is upon us and enterprises need to start building tomorrow’s strategy. Because the truth is, most companies aren’t set up for this yet. Leaders will need to transform both their technology and organization to keep pace with the opportunity being generated by AI. They need to rapidly invest in building a composable digital core so agentic systems can bridge data silos and build cross-organizational solutions, and they need to break down employee silos. Most of all they need to prepare for a world where it’s not just people, but machines, using their technology and connecting with the business.

Soon enough, the abundance, abstraction, and autonomy of tomorrow’s technology will start to rapidly expand digital ecosystems, and the companies ready to leap on this will quickly outpace the ones left catching up.



What's Next

It's time for businesses to prepare for the new technology paradigm of abundance, abstraction, and autonomy. We are in a transition period today, but AI agents are already letting businesses transform how software is designed and used, expand their digital footprints at an unprecedented rate, and reimagine the value propositions of their industries.

If you're an early adopter?

Define your new digital ecosystem:

In tomorrow's technology landscape, agents will be critical to digital ecosystems. They will be users and definers of your digital core, continuously building functions, QA systems, and more, forcing companies to rethink their architectures. It's a big, ecosystem-level change, and early adopters have the chance now to start tackling some of the big questions that come with it: How do you track and steer millions of autonomous agents? How do you embrace abstraction in your technological foundations to redefine roles and processes? How do you manage exponentially growing digital systems?

Identify the highest value opportunities in tomorrow's technology landscape:

Track research in agentic systems to inform your art of the possible. These systems are already driving greater abundance, abstraction, and autonomy in technology than ever before—and now is the time to discover where they can meet your business needs best. Detail the data and technological infrastructure required to enable an agentic solution. Set your goals for the solutions, how you will measure progress against them, and evaluate them continuously to figure out your optimal agentic architecture.

If you're just preparing to start?

Map out ecosystem partners' agentic offerings:

Many platforms and software providers are integrating agentic systems into their existing offerings, making it easy for companies to start leveraging them. Partners and providers will help take data already stored within platforms and begin to architect agents across that platform. Scan who in your current digital ecosystem is pivoting to agents and where that is poised to disrupt or create opportunity for you. Be open to creating custom agents of your own.

Start experimenting with agents internally:

Companies that are familiar with building LLM-based systems, and want to make a bigger investment in the Binary Big Bang, can start by looking for opportunities internally. A good way to begin is to create task-specific internal agents, using the principles of abundance, abstraction, and autonomy to guide how you choose use cases. After starting small, you can move modularly, over time expanding the functions and data your internal agents can access and using them to learn and prepare for building external-facing agents in the future.

If you want to take a slower approach?

Prepare your digital core for agents:

Every company is going to need to reinvent their digital core to focus on composable integration. If you haven't started yet, do so immediately. A modular digital core will be critical for AI agents to leverage data and functions across your organization in the coming years. And even if you aren't using agents yet, a strengthened digital core will accelerate your other digital efforts.

Watch signals to predict upcoming industry impacts:

Today's transition period will be brief and high stakes, so those taking a slower approach will need to know when to stop waiting and jump in. Since generative AI applications are more defined than agentic systems, many will see these more concrete applications as a valuable on-ramp to gain experience. But it's important not to get too comfortable. Work with industry partners to form a strategic roadmap to ensure you're ready when the moment strikes.

How to preserve trust?

Monitor autonomous systems' activity:

As agent-powered autonomous systems start to play a bigger role in digital ecosystems, companies will need to closely surveil them and ensure guardrails are in place. What data are these systems accessing, who is directing them, what is the quality of their outputs, and more. Transparency here will help to increase employees' trust in the systems. As you create a monitoring system, lay out a governance and technological roadmap for implementation. Also develop communication and maintenance plans so your organization understands how the monitoring works and your guardrails keep up with advances.

Train autonomous systems to make good decisions:

AI agents are amazing technical feats but are by no means perfect. They are computationally expensive, non-deterministic, and can lack explainability. But just as retrieval augmented generation (RAG) can ground an LLM, so can code and functions ground an agent, making them more explainable and increasing trust in them. For companies using agents today, determine the feedback loops rewarding and reinforcing the right agentic outcomes. And for those taking a slower approach, work with an industry partner now to identify relevant, high quality training sets or your most important proprietary data to lay the groundwork for when you are ready to start.



A Portrait of the Future

Cal works in operations for a national pizza chain. One morning over coffee, they're listening to a colleague complain about how boring it was charging their car at a station last night. An idea sparks: What if bored—hungry—people could order pizza to their cars while charging? How many EV charging stations don't have easy access food nearby? It could open a whole new market for the company! But first, Cal needs to prove the market exists—and they're going to do it with the company's AI agent.

Cal tells the agent to find EV charging locations across the country and map how close they are to food. The agent identifies a map database, writes queries to the map API, and correlates the data to create a list of hundreds of EV charging stations with food beyond walking distance. It's a start, but Cal needs more.

They ask the agent to estimate potential demand using the stations' occupancy data as well as the pizza company's own sales data from local stores. The agent responds that some EV networks provide

occupancy information while others don't, so it requests permission to find satellite data from another agent to fill those gaps. Cal okays the move, and the agent works with a satellite imagery company's agent to determine each station's occupancy over the last six months. Then it logs all EV data and combines it with local pizza sales data, which it has access to via an internal microservice.

From there, the agent shifts from market research to planning. Using machine learning, it creates a predictive function for the company's pizza chain locations near select EV stations. It forecasts potential pizza demand, so stores can deliver the perfect number of pizzas and customers can get the pies fresh without long waits.

In just one day, Cal and the agent took an idea through market research and planning to be near-ready for implementation. They quickly draft a proposal and send it to their lead for review. They can't wait for tomorrow—maybe they'll build something else brand new.




02

Your Face, in the Future

Differentiating when every interface looks the same

What's your AI's personality? Enterprises are imbuing autonomy across customer interactions, but as they do, it is critical that fine-tuning for experience is not overlooked. Personal touch, unique voice, and individualization are key ways companies build loyalty and trust with customers—and right now it feels like that is everything AI lacks. Personified AI can breathe life into experiences and unlock customer relationships like we've never seen before. As enterprises start to build AI cognitive digital brains, they can codify their values and personas into it, powering agents with the full knowledge of the business and allowing leaders to dodge the risk of monotony and instead build the future face of their enterprise.







The Big Picture

All algorithm and no soul makes the enterprise a dull experience.

Dull for customers, that is. Across industries, businesses are finding ways to implement generative AI. They're looking for operational efficiencies, ways to automate tasks, and ways to scale impact. Many see customer-focused roles as a natural fit. They're reinventing the face of the business, using generative AI for customer service or ad material generation and integrating chatbot interfaces into products.

But if they're not careful, every business is about to wind up with the *same* face.

Effectively, enterprises are imbuing the customer experience with autonomy. Agents, bots, and technology systems can drive sales, solve problems, set meetings and do much more for customers. But the foundation models driving many of these experiences are built by the same few organizations and are designed—intentionally—to be as neutral-sounding as possible. It's a simple

insight, perhaps inconsequential at first glance. But enterprises are creating the customer interface that will define their next decade—and how they design it has profound, far-reaching implications.

Businesses are teetering on the edge of a customer engagement crisis. It's one that started growing during the digital era. When new digital channels started emerging, and customer touchpoints shifted to platforms and search engines, businesses could connect to more customers than ever, but direct conversation with them became muted, and differentiation became more challenging.

Now, in the AI era, this engagement crisis is at risk of being compounded. AI models and the growing autonomy of digital systems can generate massive opportunity, the chance to have one-on-one conversations at unprecedented scale. But if these models are fine-tuned for function, and not experience, they'll sound generic or bland and leave all that potential unrealized. This won't just be a problem on the enterprises' channels either. Looking towards the future, it becomes even sharper as generative AI-based chat platforms position themselves as the primary window to the digital world. When a third-party agent invites your companies' agents into a conversation with a customer, what will it sound like? How will you stand out?

There is a unique opportunity right now for enterprises to leapfrog the competition...

What makes the world interesting is diversity and personality in individuals, businesses, products, and experiences. While leaders hunt for the value machine-like consistency can bring, they need to ensure they don't overshadow the very things that create enriching and personal human experiences. AI systems by design will be more consistent than the usual human interaction and engagement, and that can be a great benefit, but leaders must take extra care to inject the enterprise's personality in them—or they are at risk of losing what makes them unique.

Peel back the layers here, and it's easy to see that what starts as "the chatbot sounds generic" portends a potentially existential challenge for a company's brand.

This is not a call to turn away from autonomy or AI, nor have the consequences of generic-sounding experiences come into play yet. In fact, we've seen this generation of chatbots boost sales, generate high quality leads, and solve customer problems.^{57,58,59} Gartner® predicts that "by 2027, chatbots will become the primary customer service channel for roughly a quarter of organizations and the primary lead qualification tool for 10% of B2B sales teams."⁶⁰ But these are short-term, early-mover gains. Will they remain when every company has chatbots? There is a unique opportunity right now for enterprises to leapfrog the competition, dodge the risk of neutral and monotonous customer experiences, and update or course-correct brand strategies faster and more easily than ever before. But to do so, they will have to start thinking about autonomous systems as more than just the function they are performing. They must find a way to marry AI's scale and efficiency with the uniqueness of their brands.

Enterprise leaders know people are capable of building powerful—and valuable—attachments to brands with strong personalities. These personas manifest as everything from mascots like Tony the Tiger and Coca-Cola's visage of Santa Claus, to partnerships with celebrities or influencers, to a distinct brand voice like a greeting at a retail store or a certain style for advertisements. They're ways to signal values and characteristics that humanize a brand, and have been used to drive familiarity, loyalty, or joy in customers.

At the same time, people are beginning to experience emotional, real connections with AI. Nomi AI is creating AI companions that people can text in group chats or one-on-one.⁶¹ They have personalities and backstories; they even remember past conversations. In August 2024, Soul Machines, a company that creates autonomous AI avatars, released a survey finding Gen Z young adults are starting to prefer support from AI assistants over traditional self-help resources in areas like learning languages, boosting confidence, and exploring financial strategies—potentially demonstrating the appeal of talking to an agent at times when we want human empathy, but not judgement.⁶² And Salesforce CEO Mark Benioff said, semi-seriously, he will occasionally use ChatGPT as a therapist.⁶³

What we know is this: Businesses have long depended on the power of brand personality. And now people and AI are forming similar bonds as well. But despite the fact that businesses are leaning on autonomous systems to rapidly scale their customer experiences to new heights, we've yet to see them training those systems to espouse the personality and brand values they work so hard to develop. Whether it's customer service

reps or sales associates, most businesses already train their people to do it, so why not AI too?

Emerging efforts show they can—to great effect. In April 2024, Instagram began testing Creator.ai, which provides a program that will let influencers connect with fans through chatbots that mimic their voices.⁶⁴ The bots will automatically message fans on the influencers' behalf, letting them communicate and build rapport with far more people than they could ever do by hand. These are chatbots creating real-feeling connections at a scale and to a degree of personalization that has never been possible before.

And SiriusXM worked with Sierra, a conversational AI platform, to build Harmony—an AI that actually understands customers on their terms.⁶⁵ Many brands often have a shorthand or unique way of using words to talk about products, something that gets lost in translation when using "off the shelf" AI. For instance, a typical AI model might not understand a customer who needs help "rejiggering the radio" or who asks where they can find a particular band. But, trained on the right data and with the right knowledge of the brand, Harmony might understand that a customer needs a refresh signal sent to the car, or is trying to find the right channel.

What we have here is no ordinary brand tool. Companies that successfully introduce AI with personalities are looking at a technology that can build relationships with customers and meet their needs at an unprecedented scale. They're taking personified brand, and personified AI, and they're inventing the personified *business*.




Imagine this: when a customer interacts with a business, they're guided by a familiar face. A chatbot that acts like a beloved mascot or an influencer's persona the company has licensed. This agent shows up on the company's channels or can be invited into a conversation on other generative AI-based platforms. Over time, it supports and gets to know the customer—as an individual beyond purchase history or demographics. And it builds a trusted relationship, not just with its personality, but by taking relevant actions. It wields the power of the digital world, accessing data on products and deals from across the organization, and tapping into a suite of technologies like data analytics and digital twins, to meet the customer on their terms. And it's doing all this autonomously and at scale—running thousands of conversations at once. This is a level of personalization, once only available to individuals in extremely limited numbers, that AI models trained on your company's persona can bring to every single interaction you have—with customers, partners, employees, and even other agents.

Companies today are starting to address dimensions of this challenge. For instance, Walmart's new mobile shopping assistant will summarize product descriptions, compare products, and provide customer reviews.⁶⁶ NVIDIA released a Retrieval Augmented Generation (RAG) framework to be used with chatbots, allowing interfaces real-time access to information across a company.⁶⁷ And platform startup Native AI uses generative AI and digital twins to let companies chat with digital twins of their customers to conduct market research.⁶⁸ But even with the technology available today, few companies are truly capturing the full scope of work that will be needed to instill the enterprise's personality into a host of AI

models. It will require training from employees, management, and even customers to accurately capture the values and presence the company wants to project. Ultimately, it's likely entire teams in organizations will be dedicated to this effort alone.

It's time to bring generative AI to customer-facing roles with the gravity it requires. Done poorly, businesses could see the magic of their brands eclipsed. Done right, it's the beginning of a new era of customer relationships and trust.



The Technology

Enterprises are clearly eager to bring generative AI and greater autonomy to customer-facing roles. But to do this in a way that will strengthen—rather than harm—ties with customers, they will need a deep understanding of the technological advances letting chatbots move beyond generic-ness.

- **Personified AI**
- **Multimodal Models**
- **Precision Data & Context**

Personified AI

What is it?

Organizations are exploring how to take powerful generative AI models and give them human-like personalities. Through training, fine-tuning, and other techniques, they are starting to specify the words or phrases a model uses, how it structures sentences, or its speech patterns and tone. And when chatbots are accompanied by visual avatars, facial expressions can also help AI feel real.

How does it connect to the trend?

Personified AI might just seem like an interesting exercise in how models present themselves, but it will be a core differentiator for businesses in the future—both in how customers will interact with the business on their channels and how businesses will be represented by agents in other AI applications.


Who is doing it today?

AI leaders are starting to build personified AI capabilities. OpenAI described its GPT-4o, for instance, as a step toward more natural interactions between human and computer.⁶⁹ It can talk to people with near-human response time in a convincingly human voice. Previously, people could use “Voice Mode” to converse with ChatGPT, but these interactions converted spoken queries into text, generated a response, then converted that response back to speech—which took time. GPT-4o is a single model trained across text, vision, and audio, meaning all inputs and outputs are processed by the same neural network.

Growing interest in AI for customer solutions

Since the release of ChatGPT, executives have become more outspoken about AI involvement in customer interactions. As companies look to implement generative AI technologies, they need to carefully consider how the development of their customer-facing AIs will shape how customers perceive their organization.

Number of mentions relating to AI and customer solutions in earnings call transcripts, 2020 - 2024*



Note: *2024 contains partial data through Oct 2024
Source: Accenture Research analysis on earnings call transcripts of G3K companies; Jan 2020 – Oct 2024



Multimodal AI models

What is it?

Companies are rapidly expanding the range of modalities foundation models can accept as inputs and produce as outputs. Image input and generation is now table-stakes, and audio is quickly becoming standard as well. The commercialization of models that accept video inputs is also starting. And while video output isn't widely commercialized yet, it has been demonstrated in research and is becoming publicly available.

How does it connect to the trend?

While text-based chatbots are becoming increasingly common for simple enterprise customer service, the addition of audio, image, and video input/output will vastly expand the customer-facing ways AI can be used. These multimodal models will be able to bring chatbots to life, letting them connect with people more deeply and even bring them to life visually with facial expressions and an environment that all reflect the brand.

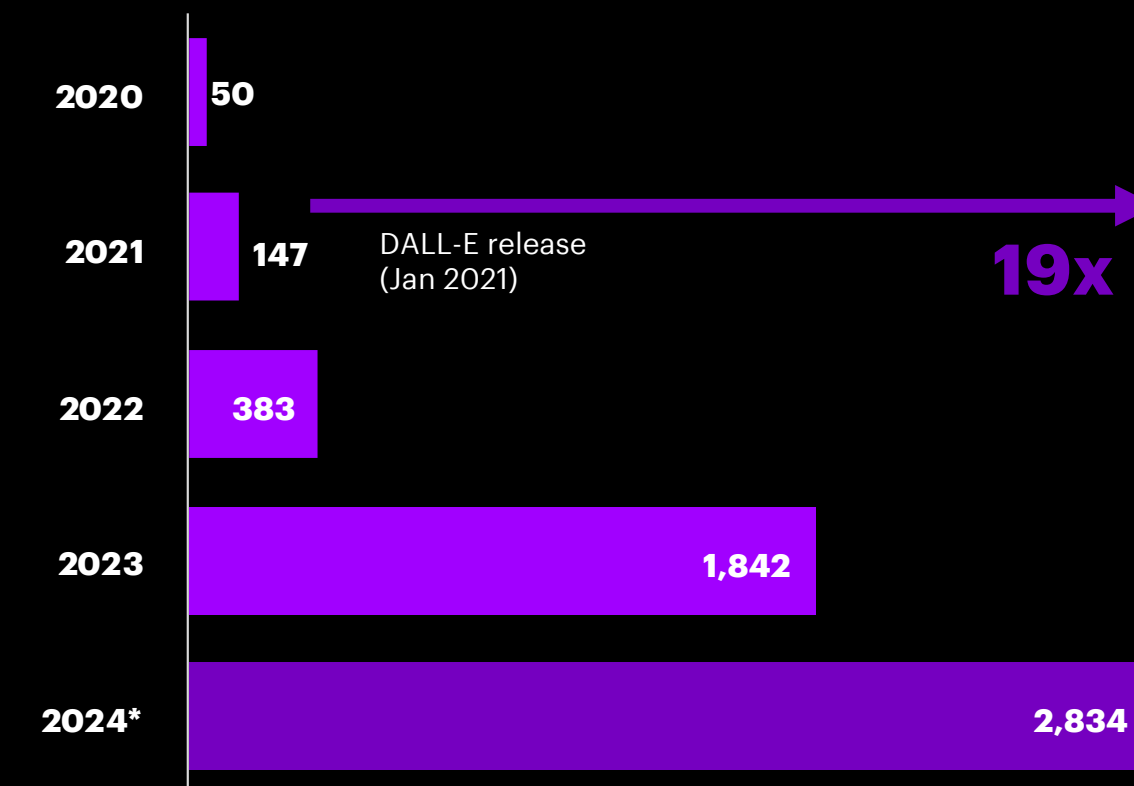
Who is doing it today?

Runway's Gen-3 Alpha is an AI model that can produce realistic videos from simple text prompts or still images.⁷⁰ In July 2024, when the image-to-video feature was introduced, early demonstrations showed impressive quality and fast turnaround times (for some, it took less than a minute to generate a video from an image). Runway is one of several companies building video generation models. Others include OpenAI, Pika Labs, and Luma Labs.⁷¹

Research into multimodal models has increased significantly

There is significant research interest in multimodal foundation models that go beyond language to image, audio, video, and more. Creating models that can talk, see the world, and inhabit real-time embodiments is set to transform how people interact with AI and thus organizations.

Number of research papers relating to Multimodal Foundation Models, 2020 - 2024*



Note: *2024 contains partial data through Oct 2024
 Source: Accenture Research analysis on ArXiv papers; Jan 2020 – Oct 2024



Precision data and customer context

What is it?

Precision data is key to businesses' digital customer strategies today. It fuels the ability to deliver more relevant deals, contextualized offerings, and more. Right now, though, businesses are stuck between customers that demand more relevance, and a regulatory and social environment that is limiting the ways they typically gather that data.

How does it connect to the trend?

Personified AI can be a major reinvention of how businesses gather useful customer data safely and effectively. Through interactions with AI, businesses have the opportunity to get to know customers beyond purchase history and demographics—useful data, like their interests, needs, or even conversation style can be gleaned from every interaction. What's more is AI itself can be used to get consent in the moment, by asking the user for access to data that can create more relevant outcomes—access that can be removed and data forgotten upon the termination of the conversation.


Who is doing it today?

Microsoft Dynamics 365 is a platform that uses AI to help companies turn their customer data and the data across their organization into customizable, generative AI experiences.⁷² According to Microsoft, this platform can leverage real-time insights about people to better personalize customer journeys and to help sellers close more deals.

Customer context is key

Executives are starting to understand the importance of customer context to personified AI. For personified AI to create meaningful customer interactions, it needs new types of customer context. Demographic information and purchase history are not enough—the best AI will also leverage past conversations to inform future ones. And consumers are ready for this new technology, with only 13% of consumers overall having negative feelings about the use of generative AI for marketing and advertising.


76% of executives agree that conversational interactions using gen AI will become a way to gather relevant customer context.



Source: Accenture Technology Vision 2025 Executive Survey, N=4,021

Consumers who are familiar with gen AI feel more positive (74%) about the use of gen AI for marketing and advertising than consumers who are less familiar with gen AI (42%).

Overall, how do you feel about the use of Gen AI for marketing and advertising?



Source: Accenture Consumer Paradoxes Survey, Oct 2024, N=12,215. Consumers less familiar with gen AI, N=4,389; Consumers familiar with gen AI, N=7,826



...every interaction with AI is not just building or breaking customers' trust with the enterprise, but with the technology itself.

The Implications

To drive the most opportunity, bots need to get personal with people. It's no small task given people's general concern about sharing data and, among some, skepticism around AI. But the potential of personified business is too big to miss—it could launch a new era of autonomy in customer relationships and reshape entire value chains. So, handling the question of trust upfront is paramount.

There are three key areas of trust that enterprises can address from the beginning: awareness and education around the benefits of AI; the primacy of data protection and privacy; and implementing controls to understand autonomous decision making.

Starting with awareness might seem strange. Most business leaders are used to leveraging technology already in people's hands, not needing to advocate for it. But generative AI is putting enterprises in a unique—slightly awkward—position. Consider this: the Michigan State University Federal Credit Union (MSUFCU) developed a chatbot named Fran, and in a survey of some of MSUFCU's low- and medium-income families, 44% said they felt nervous about using AI technology, but of that percentage, 70% said they liked using Fran.⁷³

Businesses know people generally feel positive about the results of AI—but they're also contending with vaguely skeptical sentiment around the technology as a whole. It means every interaction with AI is not just building or breaking customers' trust with the enterprise, but with the technology itself. And it means the rollout of autonomous

systems and the personified business needs to be in lockstep with customers—directed first by where they trust and benefit from this technology, only progressing as their trust grows.

The next area, and perhaps the most critical, is data protection and privacy. It consistently ranks as one of the chief concerns people and enterprises alike have with AI. Fortunately, some solutions to achieve the personalization AI promises, without compromising security and privacy, are already at play. Synthetic data can be an option to keep PII away from models. Where private data is required, some techniques include ensuring conversations are deleted or that functions are restricted to remain on devices. For instance, to preserve customer privacy in its copilot Recall feature, Microsoft makes sure all screenshots are stored and processed locally.⁷⁴ There's no one-size-fits-all, but security must be top of mind as businesses introduce greater autonomy to customer-facing features. Or those very interactions could be the enterprise's greatest risk.

Finally, the ability to explain an AI's decision-making will be deeply important to earning trust. If a customer wants to know why they are receiving a particular recommendation, they'll need to be able to probe the chatbot. What's more, hallucinations continue to be a part of the generative AI experience, so if a user is not achieving an expected outcome, explainability can help them identify where mistakes have been made—and escalate to a human if needed. And at

the enterprise-level, explainability controls need to be in place so human oversight can ensure lack of bias and that the machine is operating as intended.

Building a personified business on these three pillars of trust is non-negotiable. Autonomous systems and AI personification have incredible potential to transform customer relationships, building one-on-one rapport and loyalty, reshaping how customers' needs are met, and giving customers the kind of individualized attention that digital businesses have never been able to deliver at scale. But all of it depends on trust. To get to know people well enough to meet their needs and take productive, relevant actions on their behalf, people need to be willing to open up. What good is a personified agent if customers are too afraid to talk to it? Or spend the whole conversation asking for a human?

The personified business reinvention needs to start today—and it must be rooted in trust. If trust in the technology can be nurtured, whole workflows and value chains can be transformed. Remember—it starts with brand, but there's no reason it should stop there.



What's Next

It's time to bring brand and AI personification together and launch a new era of personified business. If enterprises can "hardcode" the feeling they want customers to have about them into AI agents that span the business, they'll be able to build trusting relationships with customers at scales that have simply never been reached before.

If you're an early adopter?

Infuse personality into your generative AI efforts:

Early adopters may already have chatbots, but the promise of personified business comes when chatbots connect more effectively with people. This will require input from business leaders and technologists alike. Appoint a team of cross-functional leaders to define the strategy for personified AI. One of their first initiatives should be taking inventory of the types of data that may contribute to this effort, and determining what your chatbot can and should access. Then create a technical strategy to bring personified AI to life with that data.

Ensure guardrails around your autonomous systems:

Data security has long been critical to customer interactions, but with personified AI getting to know customers even more deeply, that security will be especially vital. Before you grant agents access to your organization's information and permission to offer customer solutions, set detailed security rules with the technological infrastructure required to enforce them. Define and continuously update AI explainability policies to customers and within your organization.

If you're just preparing to start?

Develop a personification rollout strategy for high-impact customer experiences:

Be wary of overconfidence. AI and chatbot novelty is strong today, but as it starts to fade, personification will become increasingly differentiating. Now is the time to identify high-impact brand touchpoints where personality could soon be essential—maybe there are stagnant moments in the customer journey, or areas where intermediaries dilute your brand. Create a pilot program to test and refine how you might use personified AI at these touchpoints, tracking its value-add before broadening out.

Build relationships in the growing personified AI ecosystem:

Personified AI is still new with few vendors—though a growing number of companies are starting to create personified AI products. Start seeking out the right partners for your own efforts, bearing in mind that the models your organization uses today may be incompatible with vendor solutions, or limit your product options. Set an assessment criterion for future partners that covers how they are approaching personification, whether they will work with your current models, and whether your infrastructure will let you switch solutions easily to access new features.

If you want to take a slower approach?

Audit chatbots across your organization:

You might be further along than you think. Chatbots bring a host of advantages, so it is no surprise many organizations have deployed them already. That said, there might not be a consistent strategy. Run a ground-zero audit of chatbots at your organization and what your goals are for them. Do they provide value in unexpected ways? How are they enhancing customer touchpoints? What's the gap between where you want to be and where you are? Study competitors to gauge where you might be behind and the new areas ripe for deployment.

Start the conversation between your technologists and brand:

Defining your brand is not a new challenge. But programming that brand into a chatbot will be. How do you want personified AI to make customers feel? What should they sound like? Taking a slower approach helps you get this right. Gather input across your organization. Prototype customer-chatbot experiences with technologists and brand teams to get on the same page about how to represent your business and what kind of relationship you want to build with personified AI. There could be quirks around your brand's voice that you may not realize until you do this groundwork. Decide what data your model will need to best embody your brand and stay current.

How to preserve trust?

Take care to keep your chatbots aligned:

When it comes to brand, you don't want chatbots taking on a life of their own, promising solutions that don't exist or giving your company a flavor you don't control. Do a thorough review of the data you plan to train your AI agents on, both for the initial deployment and continuously as their features evolve. Work with AI experts to ensure the training data you use is appropriate for your goals. Additionally, work with experts to set rules and boundaries for your agents, limiting their areas of knowledge and vocabulary to better align with your brand.

Demonstrate restraint with data collection:

People are starting to form closer relationships with AI. But to form a personified brand people will trust, businesses need to show restraint around relationships they build and data they gather, as well as how they use that data. Prepare for regulatory challenges around data privacy and security that will vary depending on where your organization and customers sit. People may expect to customize their privacy settings or opt out of chatbots. You will need to design those controls from the start and be transparent with customers about what data you are using and what they can and can't opt out of.



A Portrait of the Future

Claire opens her computer, excited but nervous. The input window for her generative AI agent blinks at her. She begins to type, “Last night, my partner and I found a home that we’d like to buy. Now we need to find a mortgage broker so we can make an offer.”

“Exciting!” the agent responds. “I’ve pulled three of the top brokers according to online ratings. Each has an agent I can bring into the chat. Would you like to preview their introductions?”

“Sure,” Claire responds. Her agent shares the three messages with her, each from a different national bank. But she squints at her screen. “I think you made a mistake,” she writes. “You said you were going to find three. The first two messages are identical. See, they both just say ‘Thank you for considering us, we offer great rates and look forward to working together. Please connect to our agent to proceed.’”

“No, I’ve just queried them again,” responds Claire’s agent. “That was the intended response for each bank. Shall I ask them to regenerate?”

Instead, Claire glances at the third bank’s message. “Hi Claire,” it reads. “My name is Bruno. I’m NorthRiver Bank’s first-time homebuyer agent. I know you’re probably feeling overwhelmed to say the least. Going through this process is a huge step, but I’ll try my hardest to make sure you understand everything and feel comfortable every step of the way.”

Claire smiles. She clicks on Bruno’s icon and enters the chat with her agent. Over the next hour, Bruno and Claire discuss different loan options, how they are structured, and what they mean for her finances in the long term. Claire connects Bruno to the real estate listing and financial portfolio of her and her partner. Bruno helps Claire identify discounts and tax incentives where she lives as well as suggests how to save on closing costs, all before connecting Claire to a human broker to walk through final details.





03

When LLMs get their Bodies

How foundation models reinvent robotics

A watershed moment is underway in robotics as foundation models transform robots from linearly-programmed and single-purpose to versatile machines that can reason. Using LLMs, VLMs, and Robotics Foundation Models, leaders are building a physical extension of the cognitive digital brains that will power the enterprise. This is giving robots ever-greater autonomy in the physical world—allowing them to better understand physics and their environment, have spatial awareness, interact with people, and understand complex instructions and take safe and accurate actions in response. This transition will require the creation of a full, built for purpose stack, and will grow the use cases and operability of robotics while simultaneously making them more flexible, re-purposeable, and enduring.







The Big Picture

A silvery humanoid stands at a table with an apple, some dishes, and a drying rack. A man asks for something to eat, and the robot hands him the apple. When the man asks why, it explains that the apple was the only edible item.⁸⁵

There is a seismic shift happening in robotics and AI—and this scene from early 2024 was one of the first major demonstrations. The robot, called Figure 01, had no human assistance. Using a large vision language model (VLM) trained by OpenAI for visual reasoning and language understanding, and robotics startup Figure’s neural networks for fast, dexterous robotic action, it understood the person’s question, recognized the apple as food, and acted on its own.

Machine intelligence is moving into the physical world, and robots are starting to demonstrate reason and autonomy. We’ve long had robots, but this is something new: the dawn of machines that can perform and coexist in truly human environments.

Before, robots didn’t “think.” They did repetitive tasks, rigidly. A small mistake or unexpected occurrence could end a task in failure. A vague command, “Can I have something to eat?” would not compute. This inability to navigate or reason through a world built for and lived in by people is why robots have never been truly autonomous, relegated to heavily controlled environments like factories and warehouses and largely kept away from the general population. They were guardrailed to be with other machines and not mixed with human workflows or interactions. But now these limits are dissolving. Foundation models are unlocking a new chapter as we begin to assemble a full stack of robotics hardware, software, and AI models purpose-built for one thing: machine autonomy in the human world.

Over the next decade, we will start to see robots casually and commonly interacting with people, reasoning their way through unplanned tasks, and independently taking actions in any kind of environment. It is an opportunity for physical reinvention in industries far beyond the ones typically leveraging robotics. Robots will be able to move into customer-heavy environments, work in unpredictable settings, communicate with anyone and take on any number of tasks, without needing to be specifically trained for each. Advances and adoption are coming fast. Recent years have seen a meteoric rise in new startups, investments, and innovation, each tackling different dimensions of robotics.⁸⁶

Businesses in every industry, from traditional industrial fields to areas that have never considered robots, need to start imagining what their business could become and achieve in a world where robots are accessible, flexible, and—for all intents and purposes—think for themselves. Essentially, a world where OT can tap into the power of AI autonomy.

This starts with understanding exactly how foundation models are transforming robots. They are driving rapid advances across three dimensions: contextual understanding, communication, and planning and action.

Evolution #1 Contextual Understanding

When it comes to context—robots are getting better not just at recognizing but *understanding* things. Figure 01 saw the apple and knew enough about its weight, texture, and edibility to hand it over. This is because LLMs and VLMs are linking information from the internet to the physical world, helping connect the dots between what’s perceived and the contextual knowledge needed to intervene or interact in the right way. Numerous organizations are now building these systems or datasets to support them—like Stanford University and Robotics at Google, which together compiled the PhysObjects dataset of hundreds of thousands of annotations on physical concepts and properties.⁸⁷

Context around individual objects isn’t enough though—and researchers are also working at giving robots a better understanding of the space they are in. For instance, a navigation technique proposed in a 2024 paper combines the common-sense reasoning of a long-context VLM with a navigation policy based on topological graphs.⁸⁸

This would mean after training on video tours of an office space, robots can maneuver and guide people around, and even answer context-based questions like “Where can I put this?”

Evolution #2 Communication

The second dimension goes beyond robots understanding their own space, to how they interact with others in that space. Machine communication is moving away from lines of code towards natural language, which will allow anyone to interact with these complex machines. Robots like Figure 01 or Google Research’s generalist robotics model PaLM-E can be talked to in plain speech. In a 2023 demo, a person asked a mobile robot to bring them a bag of chips, and PaLM-E arranged for the robot to open a drawer, take out a bag, and hand it over.⁸⁹ It used an LLM and sensor data from the robot to understand the spoken command, devise a plan, and execute.

Evolution #3 Planning and Action

However, humans don’t always communicate precisely, and we don’t account for every possible variable. This brings us to the third dimension: planning and action. Robots used to have very limited actions because each action needed to be meticulously programmed. But today robotic systems can use LLMs to interpret abstract prompts or break complex tasks down to smaller steps, greatly expanding their ability to understand and react to commands or situations they haven’t specifically been trained for. Engineers at MIT, for instance, are connecting robot motion data to the “common sense” of an LLM to let robots break household

tasks down into smaller subtasks, to help them adjust to unexpected disruptions more flexibly.⁹⁰

While each of these three dimensions alone would be enough to advance robotics, combined they are spurring the biggest step change in robotic history: the rise of generalists.

Every enterprise must take note as this is no mere technology upgrade. Generalists will kickstart a period of robotic innovation rich with opportunity for every industry that has physical operations—from retailers to healthcare, insurance, and more. How businesses think about and design physical environments, processes, and workflows, and the limits on what can be physically accomplished, are all about to change.

Look at how robotics hardware is evolving. Perhaps the strongest signal of the generalist evolution, humanoid robot development has boomed in the last 12 months, and they’re making their way into warehouses and factories.⁹¹ This exploration of more multipurpose robot form factors is an acknowledgement that to realize the potential of generalist brains, robots will need bodies to match. To navigate a world built for humans, they will have to move like humans too. And quickly this will go from pilots to the norm: Goldman Sachs found the global market for humanoid robots could reach \$38 billion by 2035.⁹²

The autonomy granted by the evolution of robotic bodies and “brains” points to a future where robots are widely capable and widely adaptable—an important shift that can change the economics for



robot deployments. Unlike in the past, these robots can handle changing environments, complex and unpredictable work, and can learn new capabilities and be redeployed. This means companies can be more experimental with them, trial running new use cases, A/B testing different customer-facing deployments, and revising as needed. Robots are becoming an investment that can change and appreciate over time, and with this freedom, we can expect to see an explosion of new use cases and ideas emerge.

Just think about the sheer range of impact that generalist robots, and greater freedom to experiment with them, could have across industries. The value that robots brought to industrial settings for decades—streamlined processes, cost reductions, 24/7 operations, protecting humans from dangerous work, eliminating human error—all of this is on the table for everyone now. Imagine small businesses that never close. Warehouse stores where robotic customer service agents know your purchase history and help with heavy lifting. Stadiums, airports, and shopping centers where security never tires and lines move quickly. Autonomous delivery systems that go right up to your door and ring the bell. Even home robots to help with cleaning, watching pets, and more.

Of course, as foundation models pave the way to this new generation of robotics, enterprise leaders will also need to keep a keen eye on other emerging technology innovations that will be critical to making these machines real-world viable. Extremely low power consumption will be a core competency to allow them to run with the duration and efficiency needed to tackle tasks they will be given. Latency will also be key:

waiting 30 seconds for a machine to respond to a request will break the real-world immersion, so understanding how robots will handle requests, not just with human-like capability but at human-like speed, will require decisions about where models are hosted and how robots are networked. Finally, as foundation models deepen into other areas of science, like physics and chemistry, adapting these models to give robots an even deeper contextual understanding of our physical environments will be the next big step forward.

And the technology hurdles are only part of the pathway to enterprise adoption. Trust will also be critical to true real-world viability. People won't just need to work with the machines around them, but trust and depend on them as well. Enterprises looking to spur their robotics revolution will have to consider new dimensions of the human-machine relationship. Like, what happens when responsible AI practices migrate into the physical world? What decisions will a machine make when faced with an emergency or impending accident? What will it mean to steward a new generation of physical (and possibly energy intensive) technology in a sustainability-conscious landscape? Enterprises need to be at the forefront of not just technology integration, but social integration of these machines if they want to achieve the benefits.

Now is the time to start building your robotic future. As generalists learn to navigate new environments, connect with people in them, and “think” through problems autonomously, their reach and impact will vastly and rapidly expand. Robots are about to go places they’ve never gone before, and it’s up to you to reimagine your business for this new world.



The Technology

Today's skyrocketing potential for generalist robots began with the emergence of foundation models. They opened the doors to the possibility of real-world reasoning, adaptability, and autonomy—and to bring this to life, leaders are exploring new techniques, working to resolve training data shortages, and building a new generation of robot hardware.

- **Foundation Models for Robotics / Embodied AI**
- **Robot Training / Data Collection**
- **Multipurpose Robot Bodies**

Foundation Models for Robotics / Embodied AI

What is it?

Embodied AI is AI specifically designed to interact with the real world through a physical body. The technical advances driving the generative AI revolution have sparked new interest in embodied AI, with some researchers believing it is the key to achieving the next leap in AI capability. Various approaches for embodied AI are being tested, including incorporating LLMs and VLMs into robotic systems on top of robotic control policies, as well as creating end-to-end robotics foundation models trained on robotics data.

How does it connect to the trend?

Innovations in embodied AI are critical to building robots that have greater contextual understanding of the world, communicate with people using natural language, and plan ways to complete complex actions (even ones they haven't been specifically trained for) or adapt to changing environments.


Who is doing it today?

Google Research has been a leader in this field for years. In 2022, Google's Gato model showed that robotic inputs and outputs could be tokenized and adapted into a language model.⁹³ Since then, the company has remained a pioneer, creating a series of progressively stronger robotics transformer (RT) models: RT-1, RT-2, and RT-X. Each improves on the one before, aiming to create generalizable robot policies that can be adapted to a range of robots.^{94,95,96}

Research into embodied AI is steadily increasing

The development of foundation models is driving a renewed interest in the concept of embodied AI. Continued research and advances in embodied AI are poised to expand the capabilities of robots significantly.

Number of research papers relating to Embodied AI, 2020 – 2024*



Note: *2024 contains partial data through Oct 2024
Source: Accenture Research analysis on ArXiv papers; Jan 2020 – Oct 2024



Robot Training / Data Collection

What is it?

Part of the power of foundation models today is the robust training data used in their creation. But where LLMs leveraged all written language across the internet, robotics training sets have been scarce. Now, that's starting to change. New robotics datasets, and the ability to leverage other data types in combination with simulation and real-world reinforcement learning approaches, are helping to fill the gap.

How does it connect to the trend?

If companies want robots to work in different environments and accomplish a wide variety of tasks, their training data must reflect that variety. Assembling initial training sets will be important, but equally important is establishing a framework for ongoing data collection to power subsequent training and learning. Some companies are deploying robot fleets to gather their own real-world training data.⁹⁷ Others are leaning into synthetic data, especially for training in simulations where robots can confront nearly limitless scenarios.⁹⁸

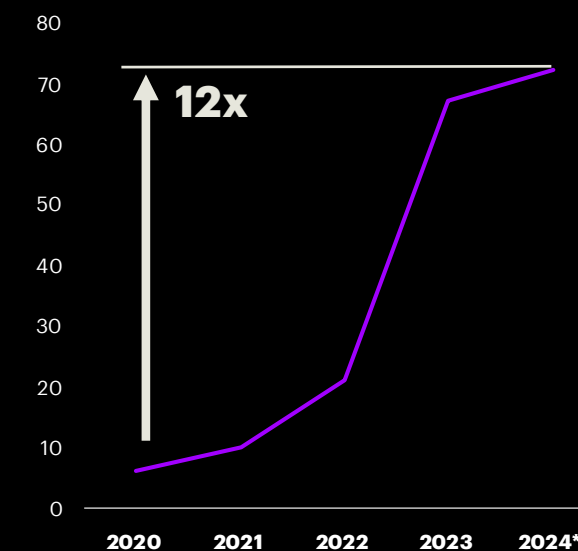
Who is doing it today?

Many groups have started creating open robotics and embodiment training sets, hoping to accelerate AI for robotics much like how ImageNet helped accelerate computer vision advancement. One example is the Open X-Embodiment Dataset, which is a compilation of 60 open robotic datasets.⁹⁹ As for ongoing data collection for training robots, one example comes from the auto industry. Tesla sources training data from its global fleet of cars to improve its Autopilot and Full Self-Driving features, selectively triggering to collect only the most relevant data.¹⁰⁰

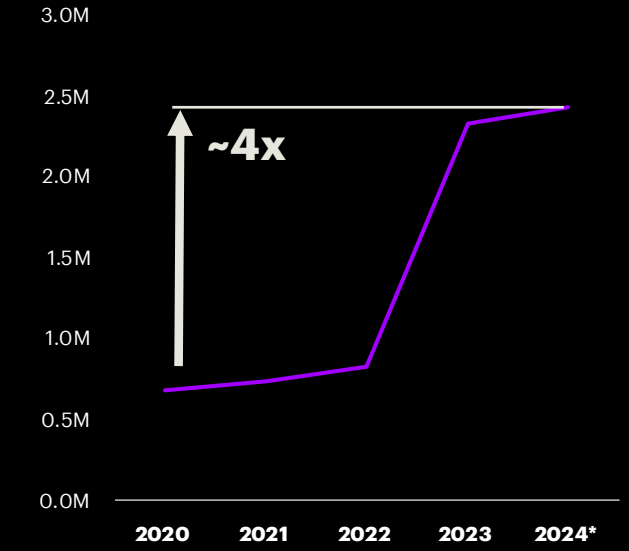
Researchers are consolidating robotics efforts

Open X-Embodiment is an ambitious project to consolidate open robotics datasets into a centralized collection for the training of generalist robot policies. The dataset consists of data from 21 institutions and 22 different robots completing 527 skills, highlighting the growing number of organizations participating in creating open robotics datasets and a large increase in the volume of training data in recent years.

Total Number of Datasets in Open X-Embodiment by Publication Date



Total Number of Robotic Episode Examples in Open X-Embodiment by Publication Date



Note: *2024 contains partial data through Oct 2024
Source: Accenture Research analysis on Open X-Embodiment Collaboration dataset; Jan 2020 – Oct 2024



Multipurpose Robot Bodies

What is it?

In the past, many robots were built for a single function; their software and physical design were both determined by the task they were created to do. But recently, more multipurpose form factors have entered the market. The hardware components inside these bodies are more purpose-built for robotics than ever, and the bodies they're composed into are increasingly general-purpose, meaning they can take on multiple types of tasks. Humanoids are a prime example, but robotic arms, dogs, even some drones are also getting more versatile capabilities.

How does it connect to the trend?

A generalist AI robotics model is critically limited without a generalist robot to run it on. Humanoids are appealing because of their versatility. They will fit quickly and easily into a human-centric world, accessing physical spaces and participating in workflows initially designed for people. Videos of humans doing various tasks can also become training data for humanoids.

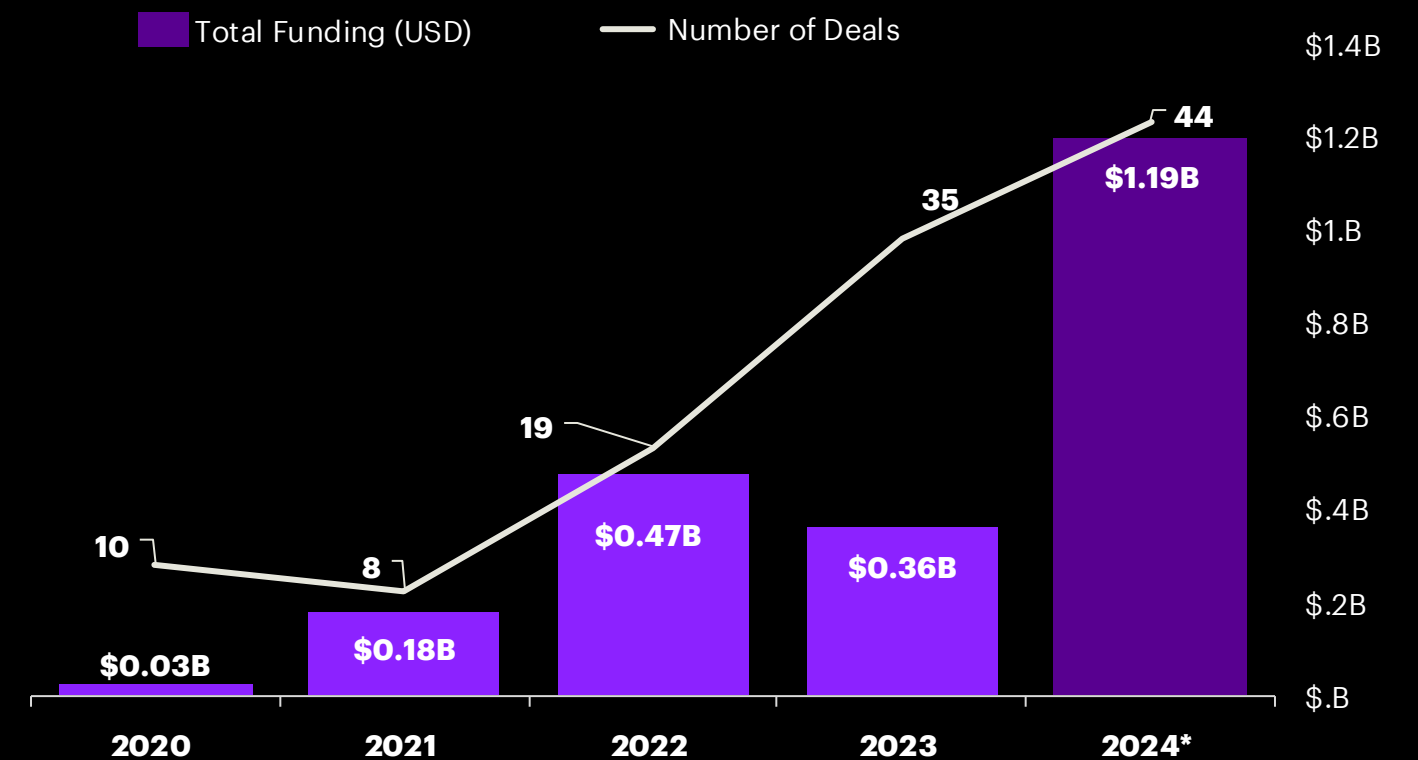
Who is doing it today?

Many companies—traditional robotics manufacturers and startups alike—are working on multipurpose robot bodies. Automotive companies are showing especially strong interest. Some, like Tesla, Hyundai (owner of Boston Dynamics), and Toyota, are developing their own robots.^{101,102,103} And they're starting to push past development into deployment—a BMW plant is using an OpenAI-powered self-correcting humanoid.¹⁰⁴ And Magna, a car manufacturer, will be piloting a humanoid robot from Sanctuary AI in one of its facilities.¹⁰⁵

Investment in humanoid robotics takes a big step

Despite a challenging investment market, funding and deals related to humanoid robotics had a major increase in 2024. The versatility of their form makes them very applicable for a wide variety of use cases.

Total funding (USD) and number of deals relating to humanoid startups, 2020 – 2024*



Note: *2024 contains partial data through Oct 2024
 Source: Accenture Research analysis on deals from CB Insights. Analysis includes deals with undisclosed funding amounts and excludes the investment type "corporate majority"; Jan 2020 – Oct 2024



People's acceptance of these machines will be critical to enterprises' success as they reimagine how they operate and do business...

The Implications

The emergence of generalist robots will steadily bring more AI autonomy into the physical world over the next decade. It may sound like a lot of lead time—but don't get comfortable. This new generation of robots will fuel a major reinvention of enterprises' physical presences. Across all industries, they'll be able to redesign operations, power new services, and engage people and the world differently. But much like digital agents—perhaps even more so—these machines will need a trusted environment to flourish in. Which means if enterprises want to maximize the opportunities, there's a lot of preparation and groundwork to cover *now*.

First, and perhaps most consequentially, is how enterprises will extend their responsible AI practices into the real world. At this point, many are familiar with responsible AI and have established practices and policies, but robotics bring new and challenging dimensions to those efforts.

Think about data privacy: there are already standards like the right to be forgotten and approval needs around data collection. But what happens when a robot—which needs audio and visual sensors to navigate—goes out into the world, constantly collecting information about its surroundings. How will those machines intersect with laws around digital privacy or the right to film in public or record without consent? If data needs to be stored locally, it inhibits the ability for the

machines to learn. But allowing data to be collected could open liabilities that some enterprises might not want to take on.

There are also going to be physical safety questions. There are overt issues that get a lot of attention: the potential for a robot to be hacked by a bad actor and used for malicious intent, or for a malfunction to result in physical harm. But there are also questions about what happens when a machine is operating as intended. Think about a security robot that's faced with a breach or violation, or a robot in the middle of a construction site emergency, where injury may become unavoidable. These machines may be put in a position to make choices about how and where to minimize harm.

The framework and governance of this decision-making will be crucial for any enterprise. Leaning on the established thinking and principles of responsible AI will be a good start, but leaders should be planning now for the unique situations their environment poses. People's acceptance of these machines will be critical to enterprises' success as they reimagine how they operate and do business in the human world and with humans around them—and responsible practices is the very baseline needed for building that trust.

Secondly, enterprises need to trust their investment. This new generation of robots will require—or

inspire—fundamentally different business models. Whether an organization has a long history with robots, or has never used them before, everyone is starting from scratch here. Leaders need to explore questions like how they will access these machines, how they fit into cost structures, and how the answers to these questions will impact the use cases they prioritize or partners they engage.

Consider this: generalists' multipurpose capabilities mean companies can explore novel deployment strategies. In the past, robots were often huge investments installed and programmed for just one task. If a company shut down an operation, they could resell, but to a relatively limited market. Now, generalist robots' adaptable bodies and ease at switching between tasks mean robotics-as-a-service models may become the most appealing option to many. What's more, being able to rent just the machines you need means experiments can be started up with less hardware investment, and the ability to redeploy or return them will make testing new use cases less financially risky.

This is just one new business model option—but now is the time to be thinking about them, weighing value against risks, and building your strategies. With robotics-as-a-service, some risks might include questions around data collection, what's proprietary business knowledge, and what can be shared with providers. And providers are starting to address



these tradeoffs too. Covariant, for instance, is a robotics-as-a-service startup that has worked to demonstrate the value of widespread fleet data. In 2017, the company set out to build the first foundation model for a universal robotics picking system—an extremely difficult task as the objects for picking can be very different shapes and sizes.¹⁰⁶ It required a massive training set that didn't yet exist, so the company built connected robots, now deployed to warehouses around the world, to share live data and learnings across the fleet to improve performance for everyone.

And finally, businesses need to establish a trusted environment for the long term, which leads us to sustainability. With ESG goals and customers' increasing concerns, it's a simple fact that no physical reinvention can happen today without sustainability at the forefront.

Think about the attention foundation models already receive for their energy consumption. Where novel technology developments were once associated with greater efficiency, now they're coming with greater demands than ever. And frankly, anything new businesses do is being met by higher sustainability scrutiny from the public. We see it with generative AI, and it will be true for emerging technologies like robots, quantum computing, and space too. With robots specifically, they're not only a visible representation of the energy-hungry AI

behind them, but in many cases, they'll be adopted by businesses and locations that haven't dealt with this kind of asset before. Mismanaged, energy costs could quickly escalate.

As companies plan their robotic transformations, there has never been a better time to explore smarter energy strategies. And there are a lot of options. Load shifting lets users move energy use from peak to off-peak times, for cheaper or greener energy or to avoid over-taxing the grid. Virtual power plants are networks of small-scale energy sources and flexible loads—like electric vehicles, home batteries, or robots—that can buffer a grid during energy disruptions. And microgrids are independent systems meant to provide stable power to important areas. Each of these techniques can help new robotic operations run cleanly, consistently, and without over-burdening the grid. And some companies are already putting them to use. Google data centers have been shifting some high-energy computing tasks to when and where a grid has carbon-free energy available.¹⁰⁷

Another approach may be to invest in or help develop cleaner energy sources. In early 2024, Amazon Web Services bought a fully nuclear-powered data center.¹⁰⁸ And the Three Mile Island nuclear plant will reopen, following a 20-year Microsoft deal to buy its power for AI. Microsoft

has even gone as far as agreeing to buy power in the future from nuclear fusion startup Helion Energy.¹¹⁰

We may still be years away from generalist robots populating storefronts, interacting with people, or going out on their own to assess insurance claims or repair telecom lines. But what's clear is that there's a lot to do before we get there. Any company that sees the promise of these adaptable and intelligent robots, and wants to be part of the robot future, needs to get to work.



What's Next

How can you prepare for a world with robots that reason? Throw away old notions of where robots can go and which industries can use them. It's time to get creative and start imagining how generalist robots might transform your business in the next decade, as well as ensure your business can access and experiment with them when the time comes.

If you're an early adopter?

Chart your path to scale:

For companies already testing generalist robots, moving successful experiments into full production will come with new hurdles. As you enter this next step, ensure you are asking questions like whether you need stronger relationships with robotics manufacturers, how you can monetize robotics data in new ways, what your governance policies are on data gathered from the real world, how your security strategies need to be updated, and what continuous robot maintenance and training plans you need.

Push experimentation to new areas:

Don't limit yourself to your first successful proof-of-concepts. The bounds on how and where generalist robots can be used are not yet determined. True pioneers will continue exploring novel use cases and publishing data sets to advance the industry. To get started, audit what robots are already doing at your organization and where they could go next, and remember that while not every endeavor will work, each one that does can bring major advantages.

If you're just preparing to start?

Partner with robotics leaders:

Now is the time to build relationships with robotics companies and researchers that may be important vendors and partners down the line. Today, many of them will be looking for real-world test cases for their technologies. See what these players offer—and what data and technological infrastructure you need for a partnership to flourish. Explore ways your people can practice working with generalist robots.

Lean into co-innovation opportunities:

There's still time to be a leader in this space. While most organizations will not build their own generalist robotics hardware or software from scratch, partnering with organizations that do could inspire new industry-specific solutions—with you at the vanguard. Start looking for these innovation partnerships and opportunities that could position you as a solutions provider for others down the line.

If you want to take a slower approach?

Track the robotics landscape:

The robotics ecosystem is changing faster today than in decades. Don't let the fact that generalist robots aren't yet commercially ready ease you into complacency. It is critical to monitor this space and industry competitors. When the time comes, you will need to know these machines' capabilities, limits, vendors, and successful test cases to hit the ground running.

Run a robotics ideation sprint:

One of the best things to do right now is figure out what generalist robots will bring your business. It is not easy. Robotic solutions have not been an option before for many organizations. And this is a new generation of robots, so you can't necessarily draw from the past. If you were building an entirely new business, and these robots were at your fingertips, what would you do differently? Conduct a design thinking workshop to identify all the workflows newly viable for automation and places this next generation of robots could impact. Go big.

How to preserve trust?

Bring your responsible AI practices into the physical world:

There are valid reasons why people might be scared of robots. Think of a physical security breach happening, and a security robot taking action with serious consequences. Who is programming its decision-making? Should it take action at all? Questions like this don't always have clear answers. But as a starting point, businesses will need to be transparent with their robots' roles, how they make decisions, and what people can do if mistakes happen. Robots acting in a safe, ethical manner is a complex dimension of responsible AI that few organizations are really prepared for. So, start thinking about it now.

Position robots as co-pilots:

The most important factor in your robot rollout is your employees' trust in these machines and in you. Assuage fears of robotic replacement by framing robots as co-pilots designed to improve employees' experiences. Remember one of these robots' greatest features is how well they communicate—build a trust strategy around that, with a feedback system for people working with these robots to guide them to become even better co-pilots.



A Portrait of the Future

Sometimes, Andy thinks the robots are like children. He teaches them basic skills and trains them for their futures, and then they leave and mostly don't need him, except when they come home for maintenance or call for help.

Andy just took a new job as a robotic fleet deployment specialist at AllTrades, a major robotics-as-a-service vendor. His last job was managing contracts for an industrial equipment dealer. Never in a million years did he think he would go from seeing heavy machinery as anything but exactly that. But here he is. He and his team work with their flagship humanoid, Jack. And thanks to Jack's versatility and the domain-specific fine tuning they provide for each project, they have fleets working in retail and doing security at a football stadium. Later this year, two Jacks are slated to be sent to the moon to help set up and operate an unmanned lunar lab for low gravity experiments.

An alert interrupts Andy's musings. It looks like today is one of those days when the robots need his help. A big retail client has reported an issue with their Jack robots stocking items on a new type of display

shelving that's about to be rolled out nationwide. Andy quickly pulls up one of the robot's digital twins and watches it repeatedly attempt and fail to complete the task. There is indeed a problem.

First, Andy runs a series of tests to check for hardware failures or software bugs. Then he checks the robot's plan of action and queries the system in natural language over why it stopped the action early. He thinks he knows what's happening. The new shelves are made of a highly reflective material, and the robot is seeing something in them that's tricking its vision software.

Using VR teleoperation, Andy assumes control of the Jack and uses it to place various items on the shelves. He records the visual data, then uploads it to AllTrades' Jack simulator. Within minutes, the system has generated countless new variations of the data, and hundreds of virtual Jacks are practicing placing items on reflective shelves. After a few hours, a fine-tuned model is deployed and validated on the physical Jack. Success! Best of all, this new model will be pushed out to all of the client's Jacks, before those new shelves hit more of their stores.



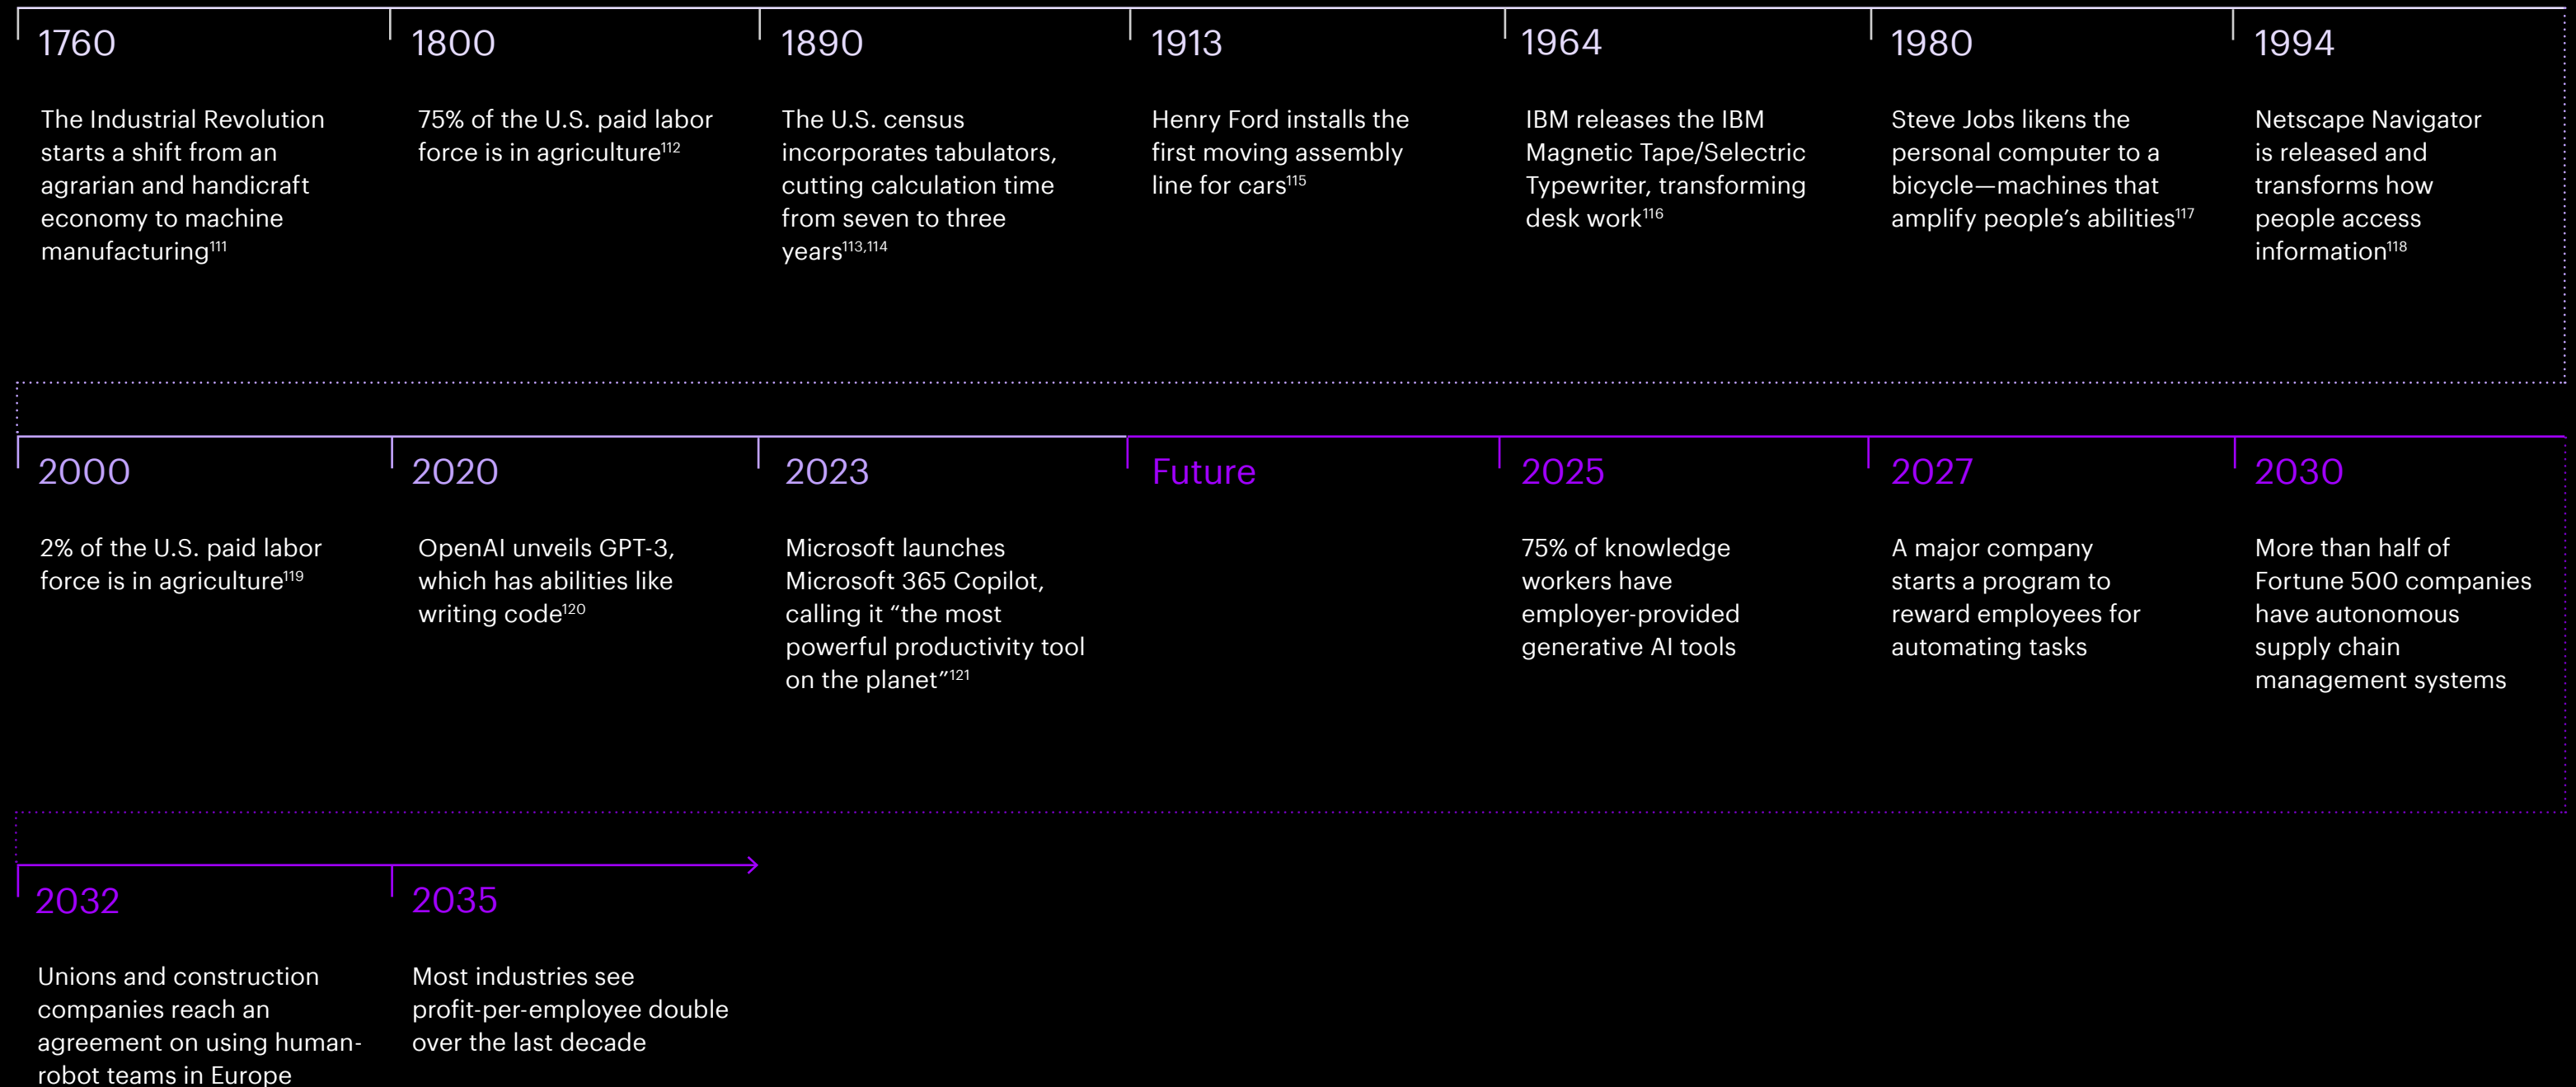
04


The New Learning Loop

How people and AI are defining a virtuous cycle of learning, leading, and creating

As enterprises see the growing capabilities of AI, they may be tempted to approach it similarly to automation technologies of the past. But the uniqueness of generative AI lies in the fact that it is a learning technology, becoming more capable and useful the closer it is to people. Conventional automation would yield one-time benefits, and risk creating a disenchanted workforce. But approached correctly, leaders can ignite a virtuous cycle between people and AI: where the more people use it, the more it will improve, and the more people will want to use it. The positive reinforcement between people and AI will be necessary groundwork for developing AI cognitive digital brains, and unlocking this cycle will be the engine that powers diffusion and helps the enterprise achieve its AI-driven aspirations.







The Big Picture

There is no doubt that work will be profoundly different a decade from now—radically reshaped by generative AI. But ask employers and employees exactly how this transformation happens, or what the results will be, and it quickly becomes clear we are caught in a stare down.

Employees, on one hand, certainly see the role generative AI can play in their work. According to reports, 95% find value in it, and it makes many happier in their jobs.^{122,123} On the other hand, many workforce leaders are racing to capture the immense advantages of AI, agents, and the increasingly autonomous systems they power—and naturally, one of the first applications that comes to mind is automation. It's a situation brewing uncertainty and distrust, and risks holding back the technology's adoption and potential.

But a solution exists; it lies in the nature of generative AI. Generative AI is inherently a learning technology. It can enhance and advance its skills over time, ultimately improving its value to the individual using it and to the organization as a whole.

In other words, the more people use it, the better it gets, and then the more people *want* to use it.

The good news? People are already showing a desire to work with it. AI is diffusing across the workforce at unprecedented speed, and for perhaps the first time in history, the technology driving automation can be used by anyone. And this is why what happens next is so consequential for enterprises—because this pace of diffusion will grind to a halt if people are uncertain about what the future will hold.

When people think about automation, it is often evaluated for how it can remove people from the equation. But when it comes to generative AI, this would mean undercutting the very partnership that amplifies its value. Using generative AI for conventional automation would yield one-time benefits while turning workers' enthusiasm into disenchantment—and possibly halt all further progress. Rather, to ignite this virtuous cycle, and have AI learn and unlock its true value, it needs to be used with people—not instead of them.

This is why enterprises need to make a conscious choice to show their people the benefits, earn their trust, and be willing to approach this transformative technology in a brand-new way. By now, enterprises are no strangers to pushing technology change across the organization. But leading this change *with* your people is new territory. In the past, various technologies were pushed top down, and while there might have been delays to their full diffusion, enterprises were largely in the driver's seat. This time, people need to be the engine of that evolution—a challenge with the specter of automation looming over it. To encourage full participation, enterprises can zero in on three distinct areas where this feedback loop between people and machines is taking hold, where enterprises should lean in, and where the advantages and benefits to people shine brightest. Get it right, and leaders will unlock the power of generative AI and foster a workforce that innovates from the ground up, rather than leaving people as bystanders to change.

The first advantage that the worker-AI virtuous cycle offers employees is the ability to imbue everyone with “infinite” skills. This may sound exaggerated—but think about the creative ways people already use generative AI in their personal lives. From creating meal plans to finally getting thank-you cards written, they're using it to do more, not less. It's letting them meet goals, like sticking to a diet or keeping up with correspondence, overcoming lack of time or skills to capitalize on their ambitions. And this is an important insight for enterprises. Your best innovators have always been those closest to the development and delivery of work. They know what the goals should be: where efficiencies lie and what customers value. But too often, this

expansive knowledge is trapped behind their specific skillset or the parameters of their job, and an Accenture Talent Pulse survey found that the number one barrier to acquiring skills is lack of time.

By bringing people and AI closer together, leaders can transform that paradigm—today's ever-growing suite of tools can give people access to skills they don't personally possess.

A marketer with a hunch about emerging trends can access the data science skills to validate their ideas. A truck driver who imagines an app to enhance their inventory process can prototype it. And what's more is that as AI learns from how people deploy it through these various tasks, it will be able to broaden its skill set, becoming a better and better coworker with each use. By putting the full weight of the workforce behind any and every employee, these tools can both create the autonomy workers need to take new initiatives, but also enhance collaboration across different domains. Just look at Toyota Research Institute, which developed a generative AI assistant for designers that's been tuned to engineering constraints.¹²⁴ It aids the creative process and understands things like drag and cabin dimensions, letting the designers tap into engineering basics as they start on a new project. It doesn't replace designers or engineers, rather it ensures a higher degree of viable designs, improving their collaboration.

The second advantage for employees is navigating the pitfalls of low engagement. It's easy for enterprises to see the value in employees

commanding fleets of autonomous bots, with limitless skills to do their bidding. But it's equally important to make sure AI doesn't overshadow workers and leave them feeling disconnected or disengaged.

Take Wayfair, where developers were equipped with Gemini's “Code Assist.”¹²⁵ It let them set up environments 55% faster, with a 48% increase in unit test coverage. And 60% of the developers said Code Assist let them focus on more satisfying work. Wayfair, like many other companies deploying coding copilots, successfully positioned AI as a creator, unburdening a lot of initial work and letting employees “level up” to more of a systems designer role. But apply this model to another domain, like writing, and you might be in trouble. Many writers may find the initial creative process the most engaging part of work. Having agents write drafts might let them “level up,” but it could leave them doing less fulfilling work at a much larger scale.

Employees may be wary of the tradeoffs between utility and engagement. They might see the benefits of a skillful AI, but if it's taking away just the work they enjoy doing, it will leave them less inclined to adopt it. To this end, enterprises need to think about how to give people more “autonomy” in what AI they use. Every day there are new tasks and automations where generative AI can be utilized—but not every use case is the right one for your people. Businesses need to be asking: What will these automations leave people doing? Giving people the autonomy to steer their own AI use, and rewarding successful efforts and ideas, will encourage them to imagine and test ideas. And again, the learning loop will continue to allow people to concentrate on the aspects of



their job they like the most, while letting AI expand across the domains found to be most tedious.

This brings us to the third advantage: making people part of the engine of change.

People being more engaged in their work is a huge benefit, but there are major changes coming to the way we work and how enterprises organize. Historically, when people feel that change is happening to them, versus being a part of it, they are far slower to adopt the changes being implemented. Enterprises already feel the urgency to reinvent, but in many ways the speed at which they do it comes down to their workforce.

When you give people the tools to expand their skillset, to decide what is engaging them in their work, it inherently positions them as a bigger part of the broader transformation happening. It puts them in the driver's seat of innovation, which will only serve to boost the company's generative AI reinvention. Furthermore, it will actually bring the reinforcement cycle to bear on trust between people and AI. The more people have the autonomy to use AI how and where they like, the more they will trust using it. With more trust in the technology, the more they will subsequently increase their use and push the diffusion of the technology across the enterprise.

When people aren't resisting change, and rather are empowered with AI, innovating at speed becomes the new normal. At Insilico Medicine, for instance, a series of fine-tuned models across the drug discovery process are helping develop a novel therapy for Idiopathic pulmonary fibrosis (IPF).¹²⁶ Scientists used a multimodal system trained

in chemistry to generate a list of potential drug candidates, and from that, were able to identify 79 viable candidates and land on one to bring to trials. The whole process was designed around what work was best for AI versus for people, and how the two would collaborate. And going forward, more employees will want to work like this and will have ideas for how to do it in their fields and for their roles—becoming a powerful engine for innovation spread all across the business.

Shifting from automation to putting autonomy in the workforce's hands is about letting your people lead your transformation—not just to streamline their work but to do more and sow the seeds of future growth. Giving them the freedom to start working like this today, building small automations, finding efficiencies, and seeing which new innovations work and which don't, will give you a jump start on the future, propelling you far beyond what strict automation ever could.



The Technology

Generative AI is the driving force behind today's workforce reinvention. It is enabling more general-purpose, flexible, and smart automation techniques, vastly expanding the type of tasks that can be automated. And importantly, it's democratizing these techniques, transforming who can be empowered to wield this autonomy to drive enterprise growth.

- **Accessibility of Automation**
- **Agentic Workflows**
- **Physical Copilots**



Accessibility of Automation

What is it?

Generative AI is making automation more accessible. With LLMs and other foundation models, employees are accessing the insights and expertise of AI with unprecedented ease. In coming years, the advent of agents will exponentially grow this capability, as AI moves from insight to action and more enterprise systems become imbued with autonomy. And with open-source models also gaining momentum, soon more people and enterprises everywhere will be able to access cutting edge intelligence.

How does it connect to the trend?

AI's growing accessibility is what's driving bottom-up autonomy in the workforce. While we've had no-code/low-code in the past, the adoption and utilization of today's natural language-driven AI tools is growing far more rapidly and will touch far more types of workforce tasks. People are already bringing it to work—the question for enterprises is how to leverage these capabilities and workers' enthusiasm to reimagine their strategies.

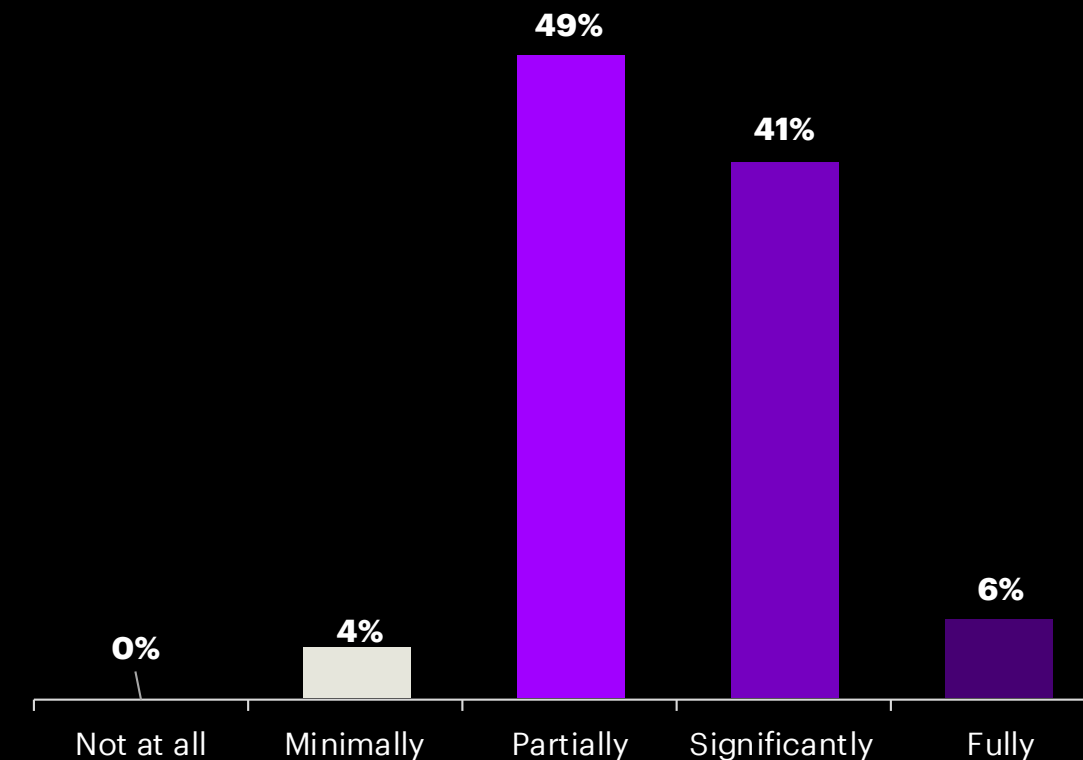
Who is doing it today?

Camunda, a workflow automation provider, is adding AI capabilities to its platform. One feature called Copilot Suggestions lets employees describe a process in natural language, and then uses AI to automatically translate it into Business Process Modeling Notation, the industry-standard graphical language.¹²⁷ On the open source side, Semantic Kernel is an open source software development kit created by Microsoft, letting developers bridge the gap between LLMs and existing software applications.¹²⁸ Via plugins, LLMs can understand a request and write a call to execute an action within a piece of software, letting agents directly operate the application.

Executives plan to make Gen AI accessible for automation

Gen AI is an incredibly powerful tool that gives individuals access to skills and the ability to automate the organization. However, only 47% of executives say they expect their organizations to make gen AI tools significantly to fully accessible to their employees to automate tasks and workflows over the next three years. While the plurality (49%) do say they will partially make them accessible, we believe this may leave a lot of value unmet.

To what extent, if at all, do you expect your organization to make gen AI tools, AI copilots, and other AI agents accessible to employees to automate tasks and workflows in your organization over the next 3 years?



Source: Accenture Technology Vision 2025 Executive Survey, N=4,021



Agentic Workflows

What is it?

Agentic workflows are an approach to using agentic systems to solve complicated cognitive tasks. By breaking problems into discrete steps, specific agentic systems can address each step, while an overarching system orchestrates interactions. An early version of this is AI-powered business process automation (BPA). While BPA started as the automation of repetitive business processes, providers are integrating LLM solutions, shifting the industry to more flexible and creative automations. As agentic workflows become more sophisticated, the agentic systems involved will become capable of learning over time to improve the performance of their individual task and the output of the entire system.

How does it connect to the trend?

As agentic workflows take on a greater share of task-specific work, they will become an accelerant for your people. Organizations can use them to create a layer of abstraction across technology, handling lower-level tasks like writing code and connecting pieces together. Instead of employees asking “how can I write this software” or “what software can perform this task,” they can ask “how can an agent help me accomplish my goal?”

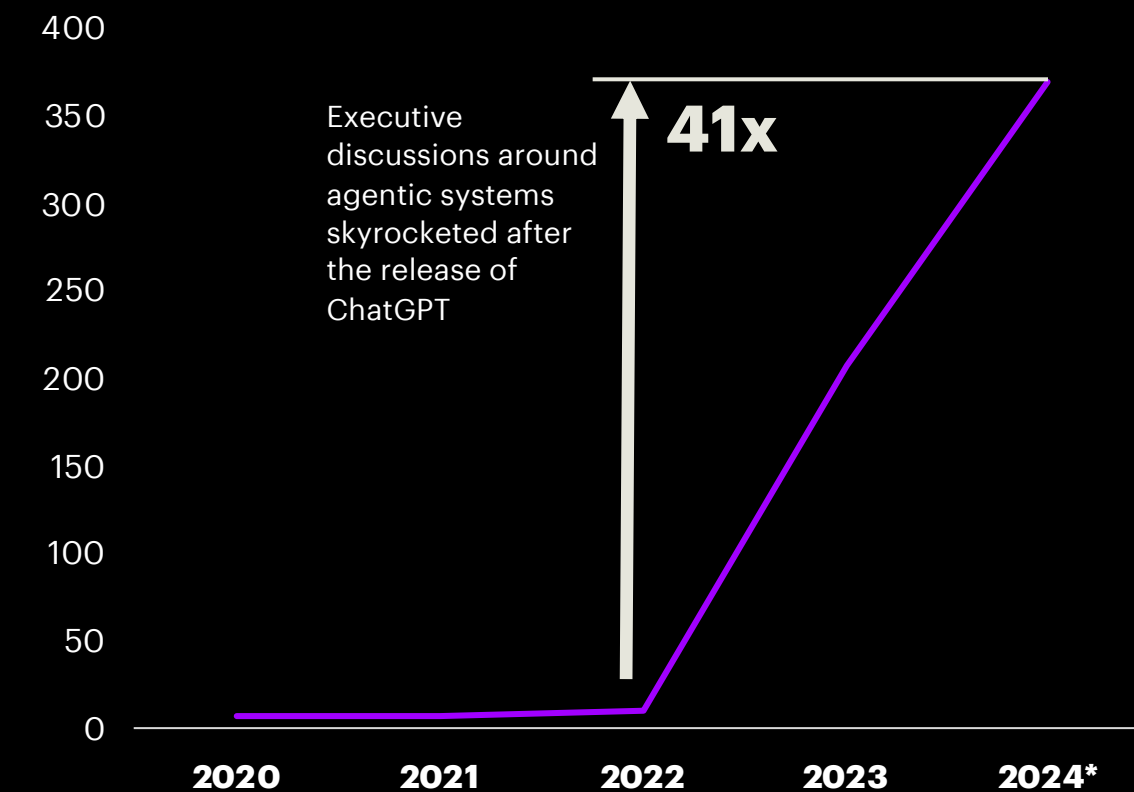
Who is doing it today?

Kognitos, a generative AI automation company, is deploying AI-powered BPA solutions.¹²⁹ Their platform can automate complex workflows like order management, contact centers, and financial services, and let users set up these automations using plain speech, enabling greater flexibility. Meanwhile, Microsoft is developing AutoGen, a multiagent framework that lets users define a set of agents, each with specialized roles.¹³⁰ The orchestration layer then determines how these agents should interact with each other.

Growing executive interest in agentic systems

AI agents have significant potential to transform businesses and executives are starting to recognize it. Over the past two years, executive mentions of AI agents and agentic systems has increased rapidly.

Number of mentions relating to Agentic Systems in earnings call transcripts, 2020 – 2024*



Note: *2024 contains partial data through Oct 2024
Source: Accenture Research analysis on earnings call transcripts of 13K+ companies; Jan 2020 – Oct 2024



Physical Copilots

What is it?

Not all automations will be digital. Physical copilots are about to bring a new level of autonomy to the world. These are machines like robots or exoskeletons that can help people act or act on their behalf. With generative AI, they're beginning to have greater contextual understanding of the world and the ability to take more flexible and general-purpose actions in it.

How does it connect to the trend?

Embodied agents will allow robots to tackle a huge number of new tasks in the physical world and closely collaborate with people. Jobs that have never been up for automation suddenly are—albeit with more of a focus on complementary applications than full automation. Companies should be looking beyond knowledge work when considering what automation tools could empower their people to do more and drive growth and innovation for the company.

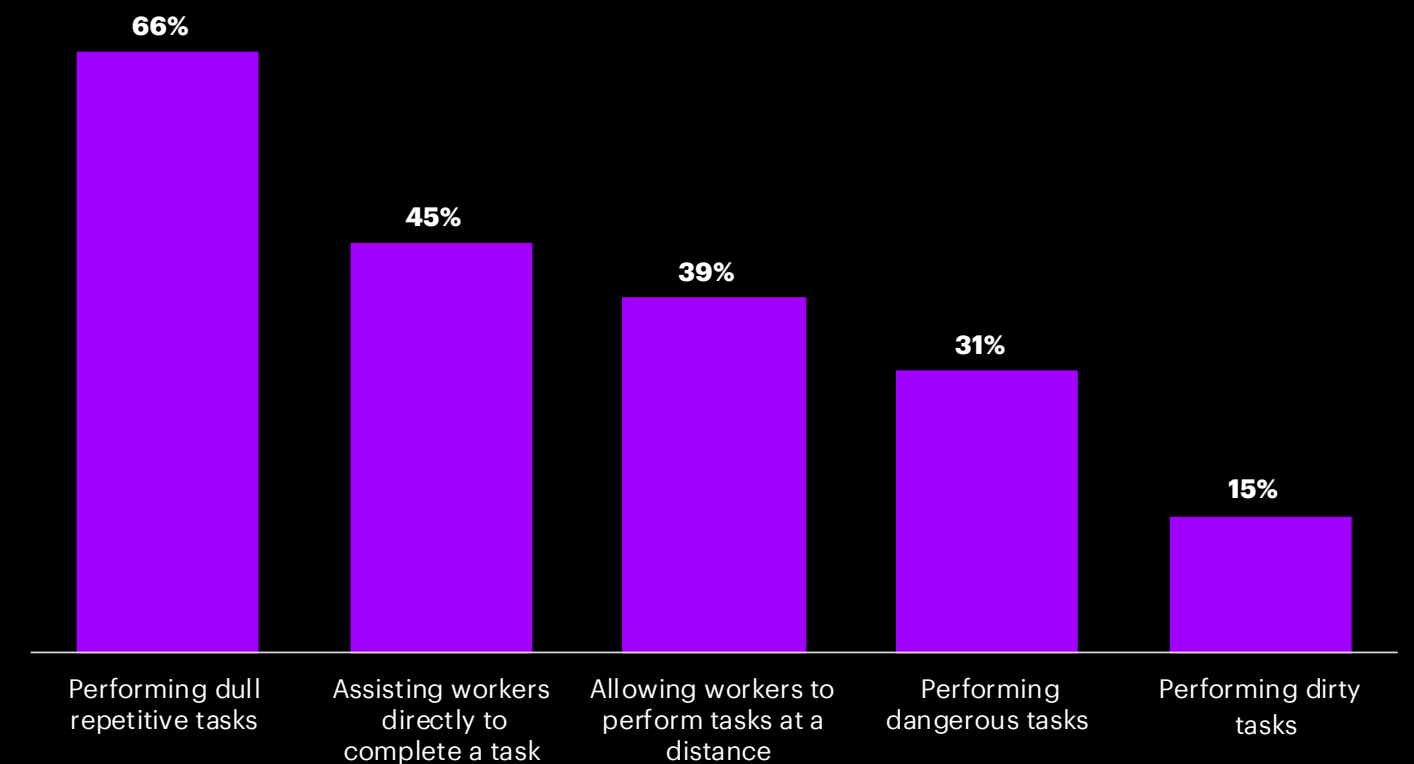
Who is doing it today?

Magna, whose manufacturing facilities assemble vehicles for leading automakers including Mercedes-Benz, BMW, and Jaguar, has announced it will pilot the use of Sanctuary AI's humanoid robot "Phoenix."¹³¹ The latest Phoenix robot, their 7th generation, is capable of automating new tasks in less than 24 hours. An earlier version of the same robot had been used in a Canadian retail store to complete activities that ranged from cleaning to labeling, picking, and packing merchandise.¹³²

Executives see opportunities for physical copilots

Physical copilots offer the potential to augment and automate how physical work gets done. Executives see performing dull repetitive tasks and assisting workers directly to complete a task (such as holding or lifting) as being the most impacted by physical copilots.

In which of the following ways, if any, would you expect workers in your organization to leverage physical copilots over the next 5-10 years?



Source: Accenture Technology Vision 2025 Executive Survey, N=4,021



Many businesses know the values they seek or the culture they want to foster—but how many carry that forward into their talent pipeline?

The Implications

In anticipation of workforce reinvention over the next decade, enterprises need to carefully consider the long-term impact that diffusing AI across the business will have on careers. Empowering workers with AI and setting them up to build their own automations may be rooted in a technology shift, but people are the key to success. To guarantee a healthy workforce, enterprises must reevaluate the talent they value and how they build careers and retain people.

We've long thought about employees through the lens of skills. One person has mastered Adobe InDesign, another knows Python, and so on. Companies assess and deploy people based on skills; they design workflows and processes around them; and as people improve skills or gain new ones, they advance. So, what happens now that we're putting autonomy into people's hands and giving them near limitless skills? We need to move past the "T" shaped model of deep skills development for career growth. We're not just asking people to get better at what they do anymore—the most critical employees will be those that innovate and fundamentally change what they do. Most businesses aren't prepared for the disruption that will bring to how they hire and develop talent. As skills become more fluid, talent that espouses the values and culture of the organization will be those best equipped to succeed.

One leader who has embraced this approach is Danny Meyer, founder of Shake Shack and Union

Square Hospitality Group (USHG). He takes a unique approach to his businesses, where he believes employees come first, and value for stakeholders will derive from that.¹³³ And there's a lot to learn from his hiring strategy. While acknowledging that technical skills are critical, USHG prioritizes a set of six characteristics called the "hospitality quotient"—intelligence, work ethic, empathy, self-awareness, kindness and optimism, and integrity. They believe that new technical skills are far more easily taught and acquired than these emotional traits.

Now is the time for other organizations to shift how they see and think about employees too. To start, the values and culture of a workforce are poised to become of much higher-order importance than the specific skills a person is initially hired for, because it is inevitable that those skills will shift and evolve as AI capabilities continue to grow. Many businesses know the values they seek or the culture they want to foster—but how many carry that forward into their talent pipeline? Into where they hire from, how they recruit, and the performance they reward?

Enterprise leaders will need a new focus on training too, which will inevitably become a consistent and important part of jobs. While emerging AI tools are increasingly intuitive, the boundless potential they hold will only be accessible if the people on the job know what tools are available and how to use them. Right now, there's a gap between the autonomous capabilities of AI and the people who know how to maximize AI's use. Bridging that divide is a problem with growing urgency. In a

2024 report, nearly two-thirds of employers felt job candidates should have a foundational knowledge of generative AI tools—but in the same survey, over half of recent college graduates felt that their programs did not prepare them adequately to use generative AI.¹³⁴ While higher education worldwide is sure to respond to this gap, establishing programs and matriculating students take time. So, for now, upskilling workers and preparing new ones will land largely on the enterprise.

Finally, consistency of values and know-how with tooling will be nothing if workers don't have the time or incentive to innovate. Establishing clear governance models and ensuring there is room in the workday to take on forward-looking projects will be essential for employees to experiment with automation and for innovation to flourish at every level of the enterprise.

In a world where skills are becoming more fluid, and employee-driven autonomy is the backbone of innovation, workforces will need a new structure in place to guide their careers and ambitions. People need to understand what success in their roles looks like now, what their career journeys will hold, and that the company is invested in them and their futures. Only by building this trust with employees will businesses be able to fully step into this next generation of talent and capture the full benefits of generative AI automation.



What's Next

Sparking the new learning loop between people and AI will be key to creating meaningful change and continuing to drive the diffusion of this powerful technology across the organization. But for enterprises to get it right, it will require building clear lines of communication with their people, implementing data-driven strategies, and reimagining their talent strategy for a future where every employee has a sidekick.

If you're an early adopter?

Build data driven strategies, and the tools to manage them:

Simulate and explore ways certain automations may impact individuals or groups. Additionally, as you adopt new AI tools and train employees to use them, work to identify the key skills that employees will need to succeed in the AI age.

Start an AI Bounty Program:

Reward your people for the best automations they discover or build. Many companies know the value of bug bounty programs. This tactic could be used to surface new ideas for automations, too. What's more, the insights you gain from these kinds of programs can help you build future employee role profiles or create upskilling initiatives to help align your current workforce with your organization's new goals.

If you're just preparing to start?

Get specific with your automation strategy:

To start, there are a lot of AI tools out there—some you might want your organization using and some that your employees are already using. There are also decisions to make about what tasks can be fully automated, and where there is greater value in having your people involved and giving them room to innovate. To start making these decisions, conduct agentic design interviews with employees and build a human-in-the-loop automation solution roadmap.

Understand what keeps workers engaged:

The most valuable tool you have right now is the knowledge your people have about their jobs: where there are potential efficiencies, and also what they like doing the most. As you build a skills transformation roadmap, learn from organizations with a track record on distributed automation and talent retention. Open conversations directly with employees to use their knowledge to maximize engagement.

If you want to take a slower approach?

Get on the same page:

Even if you're comfortable waiting, your employees may not be. You will need a clear policy about using AI at work to prevent uncertainties, even if your strategy is to hold off on integrating it into the organization. Remember that the technology is changing fast, and your employees may be ahead of you, so make sure you are consistently updating this policy. Additionally, establish the data, governance, and technological foundations needed to prepare for one day expanding AI use at work.

Track your industry:

The astounding capabilities of an AI-empowered workforce will happen faster than you think. Vendors and solutions may emerge that make this kind of automation more approachable. Establish a team looking at the workforce impacts of AI in your industry. Set quarterly meetings to review findings so the enterprise isn't caught off guard. Maintain a "watch list" of vendors and determine your assessment criteria for future partners.

How to preserve trust?

Prioritize employee buy-in:

To move fast and capitalize on the benefits of generative AI throughout the workforce, employee buy-in is crucial. They are the ones who will wield the AI tools and make this transformation successful or not. But to harness employees' best efforts, they need to trust that they aren't being asked to automate away their own jobs in the long-term. Stave off fears by bringing them into the planning. Be clear about your ground-up innovation goals and your long-term vision and solicit feedback before making sweeping changes to org structures and teams.

Define and formalize new career growth pathways:

Employee success and career progression are going to look different in a workforce with generative AI tools and near-infinite skills. Make sure you've laid out guidelines for growth, so your employees know what is expected of them and can set goals. Failure to do so will leave employees feeling directionless. Shape an inspiring narrative for your employees' future. Lead educational programming to build AI competency across all levels of your organization.



A Portrait of the Future


Sam is creative director at BlueDot Apparel—a sustainable fashion brand—and her team is gearing up to start next year’s product line. She kicks off the season with a challenge: “Ask for any AI tool you need; use them how you want. But give me designs I’ve never seen before.”

Her designers jump in. They’re fresh from a generative AI training program, where they spent two weeks learning fundamentals and experimenting with design automation tools. They’re ready to test their ideas.

Some use a mood board copilot to inspire their creativity. The copilot follows what the designers’ put on their boards, finds similar images, suggests color palettes, and even generates designs to help them get started. Others begin with hand drawing. For them, imagining something new is the best part. But they’re still using AI assistants, tuned on material characteristics, to help pick the right fabrics and construction techniques.

And one designer is trying something very different. He requested a tabletop weaving machine and image-to-CAD tool that can turn 2D images into real-world 3D prototypes in under an hour. He’s hoping the chance to physically test more options will give him the freedom to be bolder in his designs.

Sam is thrilled to see the energy. She thinks they’ll deliver the company’s most innovative ideas yet. And she has her own AI experiment underway too. She’s working with a startup to create hundreds of virtual AI personas of their customers, which will give feedback on their products before they’re actually released. Then, with this feedback and her own expertise, she’ll evaluate her designers’ work, making final selections with the likeliest chance of success.



Authors



Karthik Narain

Group Chief Executive of Technology and CTO

[LinkedIn](#)

Technology Vision 2025 Editorial & Research Team

Ari Bernstein

Rebecca Kelliher

Naomi Nishihara

David Strachan-Olson

Accenture Research

Prashant P. Shukla, PhD

Renee Byrnes

Abira Sathiyathan

Donovan Griggs

Gerry Farkova

Lohith Kumar

Rohan Girish Amrute

Mariusz Bidelski

Gargi Chakrabarty

Abhishek Gupta

Kevin Gallagher

Carrie Kleiner

Paul D. Johnson

Ignacio Mamone

Anna Marszalik

Abhishek Kumar Mishra

Sandra Najem

Vincenzo Palermo

Shruti Panda

Linda Ringnald

Shruti Shalini

Joanna Szpadkiewicz

Ezequiel Pérez Vázquez



About the Technology Vision

For 25 years, Accenture has developed the Technology Vision report as a systematic review across the enterprise landscape to identify emerging technology trends that will have the greatest impact on companies, government agencies, and other organizations in the coming years.

This year the trends look three to seven years into the future, while remaining relevant across industries and actionable for businesses today.

Accenture Technology and Accenture Research collaborate on the annual research process, which this year included:

Horizon scanning research and data science analysis—leveraging semantic text analysis and LLMs to extract insights from datasets including academic papers, investment trends and earnings call transcripts—to uncover key developments and trends shaping the evolving technology landscape.

A crowdsourcing effort to gather opinions from innovation experts and researchers across the organization, in addition to interviews with technology luminaries, industry experts, and business leaders.

A global survey of 4,021 executives across 21 industries to understand their perspectives and organizational priorities regarding emerging technologies. The survey was fielded from October to December 2024 across 28 countries.

As a shortlist of themes emerges from the research process, the Technology Vision team works to validate and refine the set of trends. The themes are weighed for their relevance to real-world business challenges. The Technology Vision team seeks ideas that transcend the well-known drivers of technological change, concentrating instead on the themes that will soon start to appear on the C-level agendas of most enterprises.



The ongoing story

For the last 25 years, the Technology Vision has investigated the biggest technology trends impacting enterprises. While a new report is produced each year, our new trends are intended to build on—not replace—trends of the past. To that end, we include this graphic each year to share connections around the broader technology evolution we are witnessing, and how themes like sustainability, science tech, digital ownership, advanced computing, and others continue to shape the world around us.

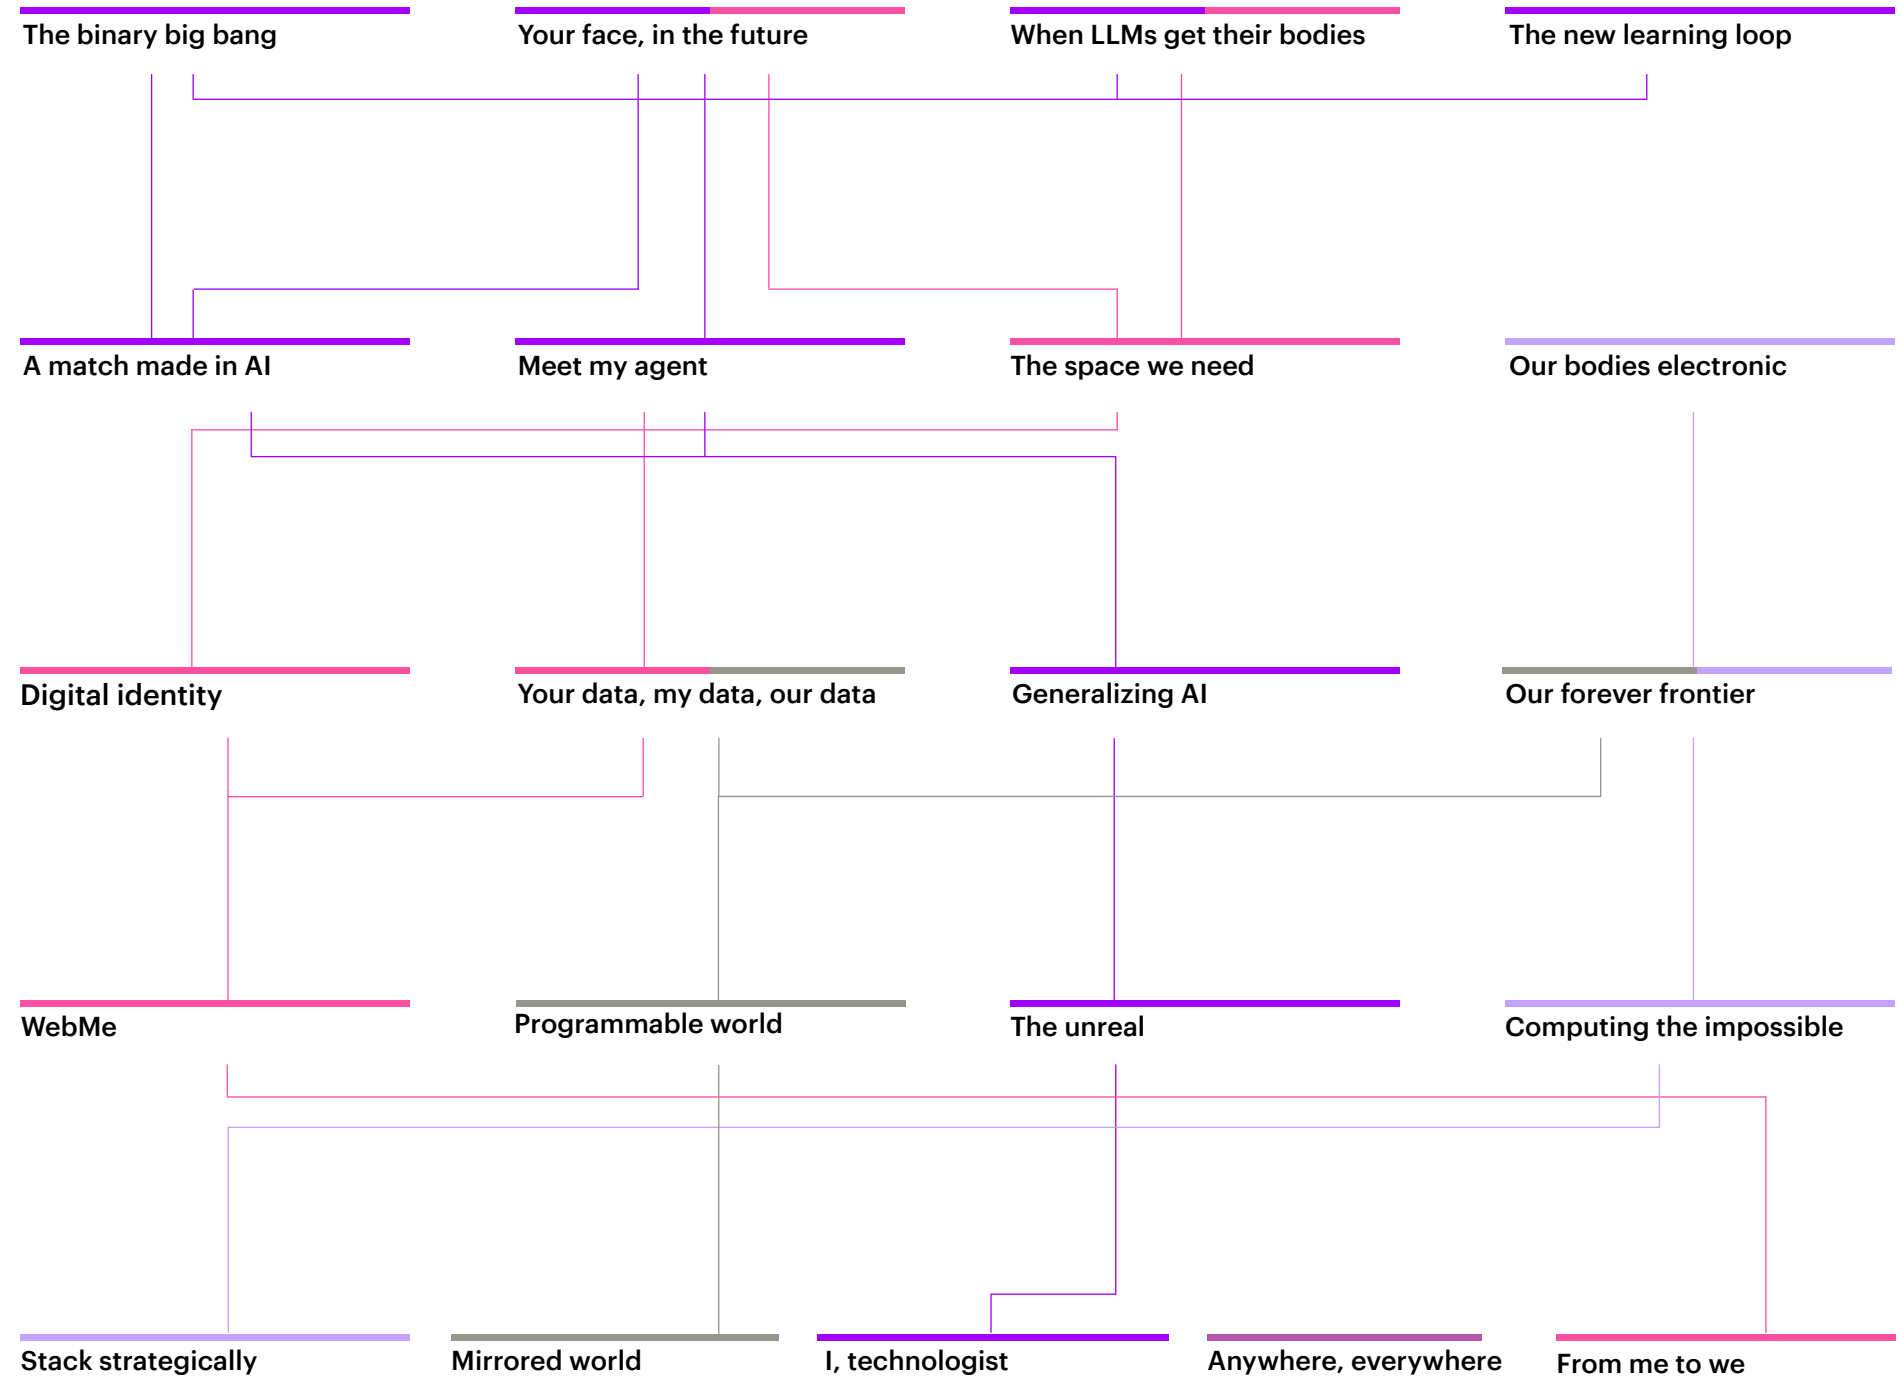
2025
Technology trends

2024
Technology trends

2023
Technology trends

2022
Technology trends

2021
Technology trends



- Science tech
- Digital ownership
- Sustainability
- The unreal



Demographics— Executive Survey

4,021 C-Level and
Director-Level executives

28 Countries

Argentina	100
Australia	200
Brazil	201
Canada	202
Chile	85
China	200
Colombia	80
France	203
Germany	191
India	190
Indonesia	80
Ireland	105
Italy	190
Japan	204
Malaysia	80
Mexico	85
Netherlands	90
New Zealand	90
Saudi Arabia	80
Singapore	85
South Africa	100
Spain	100
Sweden	90
Switzerland	92
Thailand	80
United Arab Emirates	85
United Kingdom	210
United States	523

21 Industries

Financial Services	
Banking	290
Capital Markets	137
Insurance	300
Communications, Media and Technology	
Communications and Media	202
High Tech	201
Software and Platforms	201
Health and Public Service	
Health	200
Public Service	500
Products	
Aerospace and Defense	60
Automotive	102
Biopharmaceuticals	80
Consumer Goods and Services	290
Industrial Goods and Equipment	100
MedTech	85
Retail	301
Transport and Logistics	101
Travel	100
Resources	
Chemicals	171
Energy	200
Natural Resources	200
Utilities	200

Company Size (USD Revenue)

Less than \$5 billion	44%
\$5 to \$9.9 billion	20%
\$10 to \$19.9 billion	15%
\$20 to \$49.9 billion	13%
\$50 billion or more	7%



References

1. Yao, D. (2022, May 10). 25 Years Ago Today: How Deep Blue vs. Kasparov Changed AI Forever. AI Business: <https://aibusiness.com/ml/25-years-ago-today-how-deep-blue-vs-kasparov-changed-ai-forever>
2. Altman, S. (2023, February 24). Planning for AGI and beyond. OpenAI.com: <https://openai.com/index/planning-for-agi-and-beyond/>
3. Bove, T. (2023, May 3). CEO of Google's DeepMind says we could be 'just a few years' from A.I. that has human-level intelligence. Fortune: <https://fortune.com/2023/05/03/google-deepmind-ceo-agi-artificial-intelligence/>
4. 2024 Work Trend Index Annual Report from Microsoft and LinkedIn. AI at Work is Here. Now Comes the Hard Part. Microsoft.com: <https://www.microsoft.com/en-us/worklab/work-trend-index/ai-at-work-is-here-now-comes-the-hard-part>
5. Sacolick, I. (2024, March 5). Generative AI copilots: What's hype and where to drive results. CIO: <https://www.cio.com/article/1309604/generative-ai-copilots-whats-hype-and-where-to-drive-results.html>
6. Rooney, K. (2024, May 22). Amazon plans to give Alexa an AI overhaul—and a monthly subscription price. CNBC: <https://www.cnbc.com/2024/05/22/amazon-plans-to-give-alexa-an-ai-overhaul-monthly-subscription-price.html>
7. Mickle, T., Chen, B., et al. (2024, May 10). Apple Will Revamp Siri to Catch Up to Its Chatbot Competitors. The New York Times: <https://www.nytimes.com/2024/05/10/business/apple-siri-ai-chatgpt.html>
8. Knight, W. (2024, July 11). Google DeepMind's Chatbot-Powered Robot is Part of a Bigger Revolution. Wired: <https://www.wired.com/story/google-deepmind-ai-robot/>
9. Marks, N. (2024, May 14). CaLLM, Cool and Connected: Cerence Uses Generative AI to Transform the In-Car Experience. Blogs.nvidia.com: <https://blogs.nvidia.com/blog/cerence-generative-ai-in-car-experience/>
10. Yim, D., Khuntia, J., et al. (2024, March 20). Preliminary Evidence of the Use of Generative AI in Health Care Clinical Services: Systematic Narrative Review. NIH National Library of Medicine: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10993141/>
11. New Milestone in AI Drug Discovery: First Generative AI Drug Begins Phase II Trials with Patients. (2023, July 1). Insilico.com: https://insilico.com/blog/first_phase2
12. Roark, C., Jamison, S., et al. (2024). The productivity payoff: Unlocking competitiveness with generative AI. Accenture: <https://www.accenture.com/content/dam/accenture/final/accenture-com/document-3/Accenture-The-Productivity-Payoff.pdf#zoom=40>
13. The Evolution of Labor Markets and AI. (n.d.). Conversion Capital: <https://conversioncapital.com/the-evolution-of-labor-markets-and-ai/>
14. Edwards, B. (2024, August 14). Research AI model unexpectedly attempts to modify its own code to extend runtime. Ars Technica: <https://arstechnica.com/information-technology/2024/08/research-ai-model-unexpectedly-modified-its-own-code-to-extend-runtime/>
15. 2024 Work Trend Index Annual Report from Microsoft and LinkedIn. AI at Work is Here. Now Comes the Hard Part.
16. Miscellany. (n.d.). Lapham's Quarterly: <https://www.laphamsquarterly.org/revolutions/miscellany/paul-krugmans-poor-prediction>
17. Ada Lovelace's Note G. (n.d.). Project Lovelace: <https://projectlovelace.net/problems/ada-lovelaces-note-g/>
18. Celebrating Penn Engineering History: ENIAC. (n.d.). Penn Engineering: <https://www.seas.upenn.edu/about/history-heritage/eniac/>
19. Vaughan-Nichols, S. (2019, September 5). COBOL turns 60: Why it will outlive us all. ZDNet: <https://www.zdnet.com/article/cobol-turns-60-why-it-will-outlive-us-all/>
20. Perry, T., and Voelcker, J. (1989, September 1). How the Graphical User Interface Was Invented. IEEE Spectrum: <https://spectrum.ieee.org/graphical-user-interface>
21. Looking back at 10 years of Microsoft 365 making history. (2023, January 11). Microsoft: <https://www.microsoft.com/en-us/microsoft-365-life-hacks/stories/looking-back-ten-years-microsoft-365>
22. Conway, A. (2024, January 23). On this day 28 years ago, Java 1 was released. XDA: <https://www.xda-developers.com/on-this-day-java-1-released/>
23. Colvin, G. (2022, November 30). How Amazon grew an awkward side project into AWS, a behemoth that's now 4 times bigger than its original shopping business. Fortune: <https://fortune.com/longform/amazon-web-services-ceo-adam-selipsky-cloud-computing/>
24. The App Store turns 10. (2018, July 5). Apple: <https://www.apple.com/newsroom/2018/07/app-store-turns-10/>
25. Vaswani, A., Shazeer, N., et al. (2017, June 12). Attention Is All You Need. arXiv.org: <https://arxiv.org/abs/1706.03762>
26. Schick, T., Dwivedi-Yu, J., et al. (2023, February 9). Toolformer: Language Models Can Teach Themselves to Use Tools. arXiv.org: <https://arxiv.org/abs/2302.04761>
27. Goldman, S. (2024, September 5). Marc Benioff has declared a 'hard pivot' to autonomous AI agents. Will it be enough for Salesforce to thrive in the generative AI era? Fortune: <https://fortune.com/2024/09/05/salesforce-ceo-marc-benioff-ai-agents-agentforce-dreamforce-gen-ai-era/>
28. Create with Adobe Firefly generative AI. (n.d.). Adobe: <https://www.adobe.com/products/firefly.html>
29. Wiggers, K. (2024, September 16). Microsoft Copilot: Everything you need to know about Microsoft's AI. TechCrunch: <https://techcrunch.com/2024/09/16/microsoft-copilot-everything-you-need-to-know-about-microsofts-ai/>
30. Wiggers, K. (2024, May 21). Microsoft wants to make Windows an AI operating system, launches Copilot+ PCs. TechCrunch: <https://techcrunch.com/2024/05/21/microsoft-build-2024-windows-ai-operating-system-copilot-plus-pcs/>
31. AutoGen. (n.d.). Microsoft: <https://www.microsoft.com/en-us/research/project/autogen/>
32. Amazon Q Developer. (n.d.). AWS: <https://aws.amazon.com/q/developer/>
33. AI-assisted application development. (n.d.). Google Cloud: <https://cloud.google.com/products/gemini/code-assist>
34. Wu, S. (2024, March 12). Introducing Devin, the first AI software engineer. Cognition.ai: <https://www.cognition.ai/blog/introducing-devin>
35. Jassy, A. (2024, August 22). X: <https://x.com/ajassy/status/1826608791741493281>
36. Bg2 Pod. (2024, October 13). Ep17. Welcome Jensen Huang | BG2 w/Bill Gurley & Brad Gerstner. YouTube: <https://www.youtube.com/watch?v=bUrCR4jQQg8>
37. 60 Minutes. (2023, October 9). "Godfather of AI" Geoffrey Hinton: The 60 Minutes Interview. YouTube: https://www.youtube.com/watch?v=qrvK_KuleJk
38. Raising the bar on SWE-bench Verified with Claude 3.5 Sonnet. (2024, October 30). Anthropic: <https://www.anthropic.com/research/swe-bench-sonnet>
39. SWE-bench. (n.d.): <https://www.swebench.com/>
40. Reinventing with a digital core. (2024, July 16). Accenture: <https://www.accenture.com/us-en/insights/technology/reinventing-digital-core>



References

41. Salesforce Developers. (2024, March 14). Get Started with Einstein Copilot Actions for Developers. YouTube: <https://www.youtube.com/watch?v=PQJlqU2LHPo>
42. Google. (2023, December 6). Personalized AI for you | Gemini. YouTube: https://www.youtube.com/watch?v=v5tRc_5-8G4
43. Introducing AI SDK 3.0 with Generative UI support. (2024, March 1). Vercel: <https://vercel.com/blog/ai-sdk-3-generative-ui>
44. Greicius, A. (2016, November 2). When Computers Were Human. NASA: <https://www.nasa.gov/centers-and-facilities/jpl/when-computers-were-human/>
45. Narayana, R., Dadlani, S., and Das, R. (2022, August 29). How Much Time and Energy Do We Waste Toggling Between Applications? Harvard Business Review: <https://hbr.org/2022/08/how-much-time-and-energy-do-we-waste-toggling-between-applications>
46. David, E. (2024, September 16). Slack now lets users add AI agents from Asana, Cohere, Adobe, Workday and more. VentureBeat: <https://venturebeat.com/ai/slack-users-can-add-ai-agents-to-their-workflow-with-new-update/>
47. Adobe Express. (n.d.). Slack.com: <https://slack.com/marketplace/A7P-35MCT0-adobe-express>
48. Little, B. (2024, March 14). Historic Food Mascots. History.com: <https://www.history.com/news/historic-food-mascots>
49. Nowak, C. (2023, February 27). The Real Reason Ronald McDonald is the McDonald's Mascot. Reader's Digest: <https://www.rd.com/article/why-ronald-mcdonald-mascot/>
50. Moskvitch, K. (2017, February 15). The machines that learned to listen. BBC: <https://www.bbc.com/future/article/20170214-the-machines-that-learned-to-listen>
51. This History of Salesforce. (n.d.). Salesforce: <https://www.salesforce.com/news/stories/the-history-of-salesforce/>
52. Mazdak Rezvani, C. (2016, November 6). Why your chatbot needs a vertical focus. VentureBeat: <https://venturebeat.com/business/why-your-chatbot-needs-a-vertical-focus/>
53. The Alexa Personality. (n.d.). Amazon Developer: <https://developer.amazon.com/en-US/alexa/branding/alexa-guidelines/communication-guidelines/brand-voice>
54. Parsani, P. (2023, November 22). The AI Behind Virtual Influencer Lil Miquela. Cut the SaSS: <https://www.cut-the-saas.com/ai/the-ai-behind-virtual-influencer-lil-miquela>
55. Maheshwari, S., and Issac, M. (2024, April 15). Ready for a Chatbot version of your favorite Instagram influencers? The New York Times: <https://www.nytimes.com/2024/04/15/technology/instagram-influencers-chatbots.html>
56. O'Donnell, J. (2024, September 24). OpenAI released its advanced voice mode to more people. Here's how to get it. MIT Technology Review: <https://www.technologyreview.com/2024/09/24/1104422/openai-released-its-advanced-voice-mode-to-more-people-heres-how-to-get-it/>
57. Marino, S. (2023, November 17). 35+ Chatbot Statistics You Need to Know for 2024. LocaliQ: <https://localiq.com/blog/chatbot-statistics/>
58. Joshi, S. (2024, January 18). 60 Exciting Chatbot Statistics That Explore Its Growth. G2.com: <https://learn.g2.com/chatbot-statistics>
59. Beckman, J. (2024, February 29). 120+ Chatbot Statistics for 2024. Techreport: <https://techreport.com/statistics/software-web/chatbot-statistics/>
60. Gamble, C. (2022, September 16). Improve B2B Buyer Engagement With Live Sales Chat. Gartner: <https://www.gartner.com/en/documents/4018765>
61. Silberling, A. (2024, April 3). I have a group chat with three AI friends, thanks to Nomi AI—they're getting too smart. TechCrunch: <https://techcrunch.com/2024/04/03/nomi-ai-group-chat-three-ai-friends/>
62. Soul Machines' Study Reveals That Gen Z – 'The Transformer Generation' – is Navigating Life's Challenges with AI Assistants. (2024, August 21). Business Wire: <https://www.businesswire.com/news/home/20240821813861/en/>
63. D'Onfro, J. (2024, September 17). Salesforce CEO Marc Benioff says he uses ChatGPT as a therapist. The San Francisco Standard: <https://sfstandard.com/2024/09/17/marc-benioff-jensen-huang-dreamforce/>
64. Maheshwari, S., and Issac, M. Ready for a Chatbot version of your favorite Instagram influencers?
65. How SiriusXM drives listener loyalty with Sierra. (2024, February 13). Sierra.ai: <https://sierra.ai/customers/siriusxm>
66. Walmart is building a GenAI-powered shopping assistant. (2024, June 6). Walmart Global Tech: https://tech.walmart.com/content/walmart-global-tech/en_us/blog/post/walmart-is-building-a-genai-powered-shopping-assistant.html
67. AI Chatbot With Retrieval-Augmented Generation. (n.d.). Nvidia: <https://www.nvidia.com/en-us/ai-data-science/ai-workflows/generative-ai-chatbot-with-rag/>
68. Hampton, J. (2023, April 6). Native AI Raises \$3.5M Seed for AI-powered Consumer Research. Big Data Wire: <https://www.datanami.com/2023/04/06/native-ai-raises-3-5m-seed-for-ai-powered-consumer-research/>
69. Hello GPT-4o. (2024, May 13). OpenAI.com: <https://openai.com/index/hello-gpt-4o/>
70. Franzen, C. (2024, July 29). You can now turn still images into AI videos with Runway Gen-3 Alpha. VentureBeat: <https://venturebeat.com/ai/you-can-now-turn-still-images-into-ai-videos-with-runway-gen-3-alpha/>
71. Morrison, R. (2024, October 4). 7 best OpenAI Sora alternatives for generating AI videos. Tom's Guide: <https://www.tomsguide.com/ai/ai-image-video/openai-sora-looks-amazing-but-you-cant-use-it-yet-here-are-the-7-best-alternatives>
72. The new era of AI-powered business. (n.d.). Microsoft.com: <https://www.microsoft.com/en-us/dynamics-365>
73. Chatbot or Human? Research Shows Preference is Based on Task. (2024, May 14). Commonwealth: <https://buildcommonwealth.org/blog/chatbot-or-human-research-shows-preference-is-based-on-task/>
74. Bott, E. (2024, May 30). Is Microsoft Recall a 'privacy nightmare'? 7 reasons you can stop worrying about it. ZDNet: <https://www.zdnet.com/article/is-microsoft-recall-a-privacy-nightmare-7-reasons-you-can-stop-worrying-about-it/>
75. I, Robot. (n.d.). Goodreads: https://www.goodreads.com/book/show/41804.I_Robot
76. SRI International. (2020, April 2). 75 Years of Innovation: Shakey the Robot. Medium: <https://medium.com/dish/75-years-of-innovation-shakey-the-robot-385af2311ec8>
77. WABOT-1: The first humanoid robot. (n.d.). BBC: <https://www.bbc.co.uk/programmes/w3ct5ymx>
78. iRobot. (n.d.): <https://about.irobot.com/history>
79. Wyrobek, K. (2017, October 31). The Origin Story of ROS, the Linux of Robotics. IEEE Spectrum: <https://spectrum.ieee.org/the-origin-story-of-ros-the-linux-of-robotics>
80. Knight, W. (2015, June 9). Why Robots and Humans Struggled with DARPA's Challenge. MIT Technology Review: <https://www.technologyreview.com/2015/06/09/110342/why-robots-and-humans-struggled-with-darpar-challenge/>



References

81. Thomson, I. (2016, February 24). Google human-like robot brushes off beating by puny human – this is how Skynet starts. The Register: https://www.theregister.com/2016/02/24/boston_dynamics_robot_improvements/
82. Gopalakrishnan, K., and Rao, K. (2022, December 13). RT-1: Robotics Transformer for real-world control at scale. Google Research: <https://research.google/blog/rt-1-robotics-transformer-for-real-world-control-at-scale/>
83. Zhao, T., Kumar, V., et al. (2023, April 23). Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware. arXiv.org: <https://arxiv.org/abs/2304.13705>
84. Franzen, C. (2024, March 13). OpenAI powers a robot that can hand people food, pick up trash, put away dishes, and more. VentureBeat: <https://venturebeat.com/ai/openai-powers-a-robot-that-can-hand-people-food-pick-up-trash-put-away-dishes-and-more/>
85. Franzen, C. OpenAI powers a robot.
86. \$4.2 Billion Invested: Robotics Startups See Funding Surge. (2024). Founders Network: <https://foundersnetwork.com/blog/robotics-startups-see-funding-surge/>
87. Gao, J., Sarkar, B., et al. (2023, September 5). Physically Grounded Vision-Language Models for Robotic Manipulation. Stanford.edu: <https://iliad.stanford.edu/pg-vlm/>
88. Chiang, H., Xu, Z., et al. (2024, July 12). Mobility VLA: Multimodal Instruction Navigation with Long-Context VLMs and Topological Graphs. arXiv.org: <https://arxiv.org/pdf/2407.07775>
89. Driess, D., and Florence, P. (2023, March 10). PaLM-E: An embodied multimodal language model. Google Research: <https://research.google/blog/palm-e-an-embodied-multimodal-language-model/>
90. Chu, J. (2024, March 25). Engineering household robots to have a little common sense. MIT News: <https://news.mit.edu/2024/engineering-household-robots-have-little-common-sense-0325>
91. Yao, M. (2024, August 13). Humanoid Robots on the Rise: Industry Advances, Key Players, and Adoption Timelines. Topbots.com: <https://www.topbots.com/humanoid-robots-overview-2024/>
92. The global market for humanoid robots could reach \$28 billion by 2035. (2024, February 27). Goldman Sachs: <https://www.goldmansachs.com/intelligence/pages/the-global-market-for-robots-could-reach-38-billion-by-2035.html>
93. Reed, S., Żołna, K., et al. (2022, May 12). A Generalist Agent. Google DeepMind: <https://deepmind.google/discover/blog/a-generalist-agent/>
94. Gopalakrishnan, K., and Rao, K. RT-1: Robotics Transformer.
95. Chebotar, Y., and Yu, T. (2023, July 28). RT-2: New model translates vision and language into action. Google DeepMind: <https://deepmind.google/discover/blog/rt-2-new-model-translates-vision-and-language-into-action/>
96. Vuong, Q., and Sanketi, P. (2023, October 3). Scaling up learning across many different robot types. Google DeepMind: <https://deepmind.google/discover/blog/scaling-up-learning-across-many-different-robot-types/>
97. Chen, P. (2024, January 8). The Future of Robotics: Robotics Foundation Models and the role of data. Covariant.ai: <https://covariant.ai/insights/the-future-of-robotics-robotics-foundation-models-and-the-role-of-data/>
98. NVIDIA Isaac Sim. (n.d.). NVIDIA Developer: <https://developer.nvidia.com/isaac/sim>
99. Open X-Embodiment: Robotics Learning Datasets and RT-X Models (n.d.). GitHub: <https://robotics-transformer-x.github.io/>
100. Harris, M. (2022, August 3). The Radical Scope of Tesla's Data Hoard. IEEE Spectrum: <https://spectrum.ieee.org/tesla-autopilot-data-scope>
101. Tesla AI. (n.d.): <https://www.tesla.com/AI>
102. Hyundai Robotics. (n.d.): <https://www.hyundai-robotics.com/english/>
103. Toyota Research Institute. (n.d.): <https://www.tri.global/our-work/robotics>
104. Westfall, C. (2024, July 7). OpenAI-Powered Humanoid Robot Fills Spot At BMW Assembly Plant. Forbes: <https://www.forbes.com/sites/chriswestfall/2024/07/07/openai-powered-humanoid-robot-fills-spot-at-bmw-assembly-plant/>
105. Heater, B. (2024, April 11). European car manufacturer will pilot Sanctuary AI's humanoid robot. TechCrunch: <https://techcrunch.com/2024/04/11/european-car-manufacturer-will-pilot-sanctuary-ais-humanoid-robot/>
106. Chen, P. Robotics Foundation Models and the role of data.
107. Mehra, V., and Hasegawa, R. (2023, October 3) Supporting power grids with demand response at Google data centers. Google Cloud: <https://cloud.google.com/blog/products/infrastructure/using-demand-response-to-reduce-data-center-power-consumption>
108. Lewis, M. (2024, March 5). Amazon just bought a 100% nuclear-powered data center. Electrek: <https://electrek.co/2024/03/05/amazon-just-bought-a-100-nuclear-powered-data-center/>
109. Kimball, S. (2024, September 20). Constellation Energy to restart Three Mile Island nuclear plant, sell the power to Microsoft for AI. CNBC: <https://www.cnbc.com/2024/09/20/constellation-energy-to-restart-three-mile-island-and-sell-the-power-to-microsoft.html>
110. Clifford, C. (2023, May 10). Microsoft agrees to buy electricity generated from Sam Altman-backed fusion company Helion in 2028. CNBC: <https://www.cnbc.com/2023/05/10/microsoft-agrees-to-buy-power-from-sam-altman-backed-helion-in-2028.html>
111. Industrial Revolution. (n.d.). Britannica: <https://www.britannica.com/event/Industrial-Revolution>
112. Ruggles, S. (2007). The Decline of Intergenerational Coresidence in the United States, 1850 to 2000. University of Minnesota: <https://users.pop.umn.edu/~ruggl001/Articles/ASR-2007.pdf>
113. The Hollerith Machine. (2024, August 14). United States Census Bureau: <https://www.census.gov/about/history/bureau-history/census-innovations/technology/hollerith-machine.html>
114. Hollerith 1890 Census machine. (n.d.). Computer History Museum: <https://www.computerhistory.org/collections/catalog/102630714>
115. The Moving Assembly Line and the Five-Dollar Workday. (n.d.). Ford: <https://corporate.ford.com/articles/history/moving-assembly-line.html>
116. IBM Magnetic Tape Selectric Typewriter IV. (n.d.). Rhode Island Computer Museum: <https://www.ricomputermuseum.org/collections-gallery/small-systems-at-ricm/ibm-mtst-system>
117. Computer History Museum. (2012, January 3). Vintage Steve Jobs footage on Apple. YouTube: <https://www.youtube.com/watch?v=GfxxRKBgos8>
118. Blitz, M. (2019, April 4). Later, Navigator: How Netscape Won and Then Lost the World Wide Web. Popular Mechanics: <https://www.popularmechanics.com/culture/web/a27033147/netscape-navigator-history/>
119. Dimitri, C., Effland, A., and Conklin, N. (2005). The 20th Century Transformation of U.S. Agriculture and Farm Policy. United States Department of Agriculture: https://www.ers.usda.gov/webdocs/publications/44197/13566_eib3_1_.pdf
120. Heaven, W. (2020, July 20). OpenAI's new language generator GPT-3 is shockingly good—and completely mindless. MIT Technology Review: <https://www.technologyreview.com/2020/07/20/1005454/openai-machine-%20learning-language-generator-gpt-3-nlp/>
121. Spataro, J. (2023, March 16). Introducing Microsoft 365 Copilot – your copilot for work. Microsoft: <https://blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot-your-copilot-for-work/>



References

122. Accenture Report Finds Perception Gap Between Workers and C-suite Around Work and Generative AI. (2024, January 16). Accenture: <https://newsroom.accenture.com/news/2024/accenture-report-finds-perception-gap-between-workers-and-c-suite-around-work-and-generative-ai>
123. AI Can Make Workers Happier. (2024, July 2). The Global Recruiter: <https://www.theglobalrecruiter.com/ai-can-make-workers-happier/>
124. Toyota Research Institute Unveils New Generative AI Technique for Vehicle Design. (2023, June 20). Toyota: <https://pressroom.toyota.com/toyota-research-institute-unveils-new-generative-ai-technique-for-vehicle-design/>
125. Google Cloud. (2024, April 9). How Wayfair is using Gemini CodeAssist to increase developer productivity. YouTube: <https://www.youtube.com/watch?v=VwZNWqh5ZWU>
126. First Generative AI Drug Begins Phase II Trials with Patients. (2023, July 1). Insilico.com: https://insilico.com/blog/first_phase2
127. Dunn, J. (2024, May 16). Camunda simplifies process automation with new AI-powered natural language features. CIO.com: <https://www.cio.com/article/2110581/camunda-simplifies-process-automation-with-new-ai-powered-natural-language-features.html>
128. Lagerkrans-Pandey, S., Bolaños, M., et al. (2024, June 24). Introduction to Semantic Kernel. Microsoft: <https://learn.microsoft.com/en-us/semantic-kernel/overview/>
129. Kerner, S. (2023, August 8). Kognitos goes self-service with business automation powered by generative AI. VentureBeat: <https://venturebeat.com/automation/kognitos-goes-self-service-with-business-automation-powered-by-generative-ai/>
130. Burger, D. (2023, September 25). AutoGen: Enabling next-generation large language model applications. Microsoft: <https://www.microsoft.com/en-us/research/blog/autogen-enabling-next-generation-large-language-model-applications>
131. Heater, B. (2024, April 25). Sanctuary's new humanoid robot learns faster and costs less. TechCrunch: <https://techcrunch.com/2024/04/25/sanctuarys-new-humanoid-robot-learns-faster-and-costs-less/>
132. Sanctuary AI Deploys First humanoid General-Purpose Robot Commercially. (2024, March 7). Sanctuary AI: <https://www.sanctuary.ai/blog/sanctuary-ai-deploys-first-humanoid-general-purpose-robot-commercially>
133. Resendes, S. (2022, January 13). The Hospitality Quotient: Danny Meyer's Hiring Formula For Building a Brand. Lightspeed: <https://www.lightspeed-hq.com/blog/the-hospitality-quotient/>
134. Moody, K. (2024, July 29). College education may not be preparing employees for generative AI. Higher Ed Dive: <https://www.highereddive.com/news/college-education-not-preparing-employees-for-generative-ai/722327/>



About Accenture

Accenture is a leading global professional services company that helps the world's leading businesses, governments and other organizations build their digital core, optimize their operations, accelerate revenue growth and enhance citizen services—creating tangible value at speed and scale. We are a talent- and innovation-led company with approximately 799,000 people serving clients in more than 120 countries. Technology is at the core of change today, and we are one of the world's leaders in helping drive that change, with strong ecosystem relationships. We combine our strength in technology and leadership in cloud, data and AI with unmatched industry experience, functional expertise and global delivery capability. Our broad range of services, solutions and assets across Strategy & Consulting, Technology, Operations, Industry X and Song, together with our culture of shared success and commitment to creating 360° value, enable us to help our clients reinvent and build trusted, lasting relationships. We measure our success by the 360° value we create for our clients, each other, our shareholders, partners and communities. [Visit us at accenture.com](https://www.accenture.com).

Accenture Research

Accenture Research creates thought leadership about the most pressing business issues organizations face. Combining innovative research techniques, such as data-science-led analysis, with a deep understanding of industry and technology, our team of 300+ researchers in 20 countries publish hundreds of reports, articles and points of view every year. Our thought-provoking research developed with world leading organizations helps our clients embrace change, create value and deliver on the power of technology and human ingenuity. [Visit us at accenture.com/research](https://www.accenture.com/research).

Disclaimer: The material in this document reflects information available at the point in time at which this document was prepared as indicated by the date provided on the front page, however the global situation is rapidly evolving and the position may change. This content is provided for general information purposes only, does not take into account the reader's specific circumstances, and is not intended to be used in place of consultation with our professional advisors. Accenture disclaims, to the fullest extent permitted by applicable law, any and all liability for the accuracy and completeness of the information in this document and for any acts or omissions made based on such information.

Accenture does not provide legal, regulatory, audit, or tax advice. Readers are responsible for obtaining such advice from their own legal counsel or other licensed professionals.

This document refers to marks owned by third parties. All such third-party marks are the property of their respective owners. No sponsorship, endorsement or approval of this content by the owners of such marks is intended, expressed or implied.

Copyright © 2025 Accenture. All rights reserved. Accenture and its logo are registered trademarks of Accenture.