

AGENTIC

The Rise of AI Agents

About this document

This paper provides an overview of the growing role of AI agents for established work functions, through definitions, types of agent, examples of existing agentic applications, future application of AI agents and their challenges and risks. A NotebookLM audio conversation version of this paper is [available at this link](#). This paper was written by James Carson.

Key section headings:

What are AI agents, and what is the 'Agentic'?	3
5 types of AI agent	6
Existing Agentic tools	7
Applications of AI agents	11
Challenges and risks	14
4 viewpoints from Big Tech	17

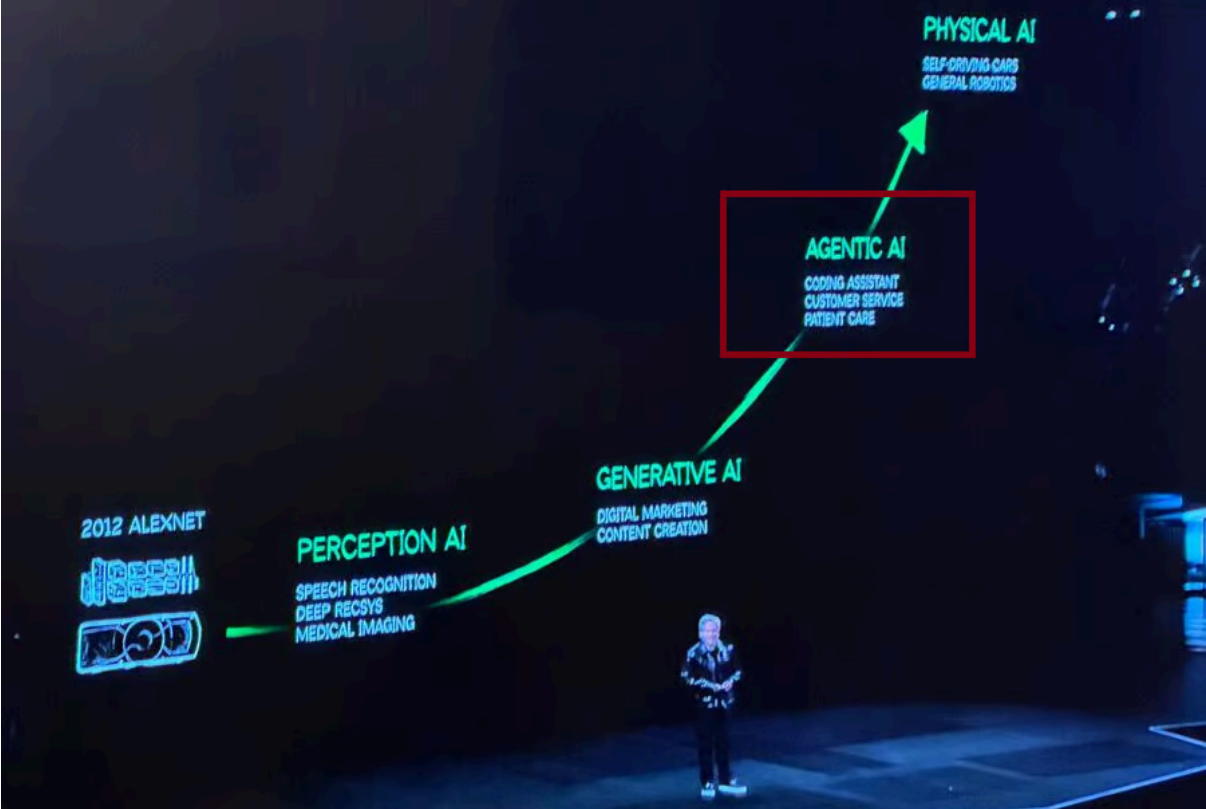
About *AGENTIC*

Agentic is an AI consultancy based in the UK, founded by James Carson, Adam Bunn and Harry Atkins in summer 2024. We work with companies with significant content output, and provide consultancy in how AI and automation can streamline this. We also work with AI companies on their content and marketing approach. We aim to distribute a white paper on LinkedIn every fortnight.

All of our white papers are available to our newsletter subscribers. Go on, [sign up](#)... it's free.

Introduction

At CES in January 2025, one of the many gigantic slides Nvidia CEO Jensen Huang stood before showed the direction of travel in Artificial Intelligence. First came Perception, illustrated through speech recognition assistants like Apple's Siri or Amazon Alexa, or image recognition such as in Apple Photos or Google Images.



Jensen Huang onstage at CES, stood before a slide showing the expected phases of AI.

Next came Generative, hallmarked by the arrival of ChatGPT in November 2022, along with a host of other AI tools across image, voice and video generation. The most difficult frontier of generative media was video, but that was mostly taken care of with a flurry of releases during 2024.

In the last three or so years, the Generative phase has seen a blend of astonishment and disappointment in equal measure. AI tools and their features get readily demo-d and hyped, leading some workers to halt their day job for a moment to try them out. However, few attempts to really unlock productivity end in immediate success, leading to relatively small adoption rates considering the excitement. We estimate that 'regular' AI adoption at work is no more than 20% (usage 3x a week or more), while the remarkable range of tools and compliance issues is confusing to many businesses. A further hurdle is the

lingering perception that generative AI just isn't quite 'good enough.' That it causes just as many problems as it solves, that what it creates looks uncanny or that it hallucinates and thus presents risk.

That said, much like the above slide suggests, we believe we're both reaching a take off moment and entering into the next 'Agentic' phase. This is not to say the Generative phase is 'completed'. We will still encounter the not 'good enough' issues for some time yet, but an oft quoted adage from Professor Ethan Mollick is that, 'This is the worst version it'll ever be.' It's impossible to argue with that.

The fundamentals of how generative AI could potentially enhance productivity - through text, image, voice or video generation, or analytical tasks concerning data or text - have all been answered, albeit with varying degrees of success. A tool or platform now exists that enables any one of these things. But there is not much yet that does a great job of melding them together very well. This is the solid base for the next phase - the 'Agentic'.

What are AI agents, and what is the 'Agentic'?

When we say 'Agentic' what we mean is a phase of AI which utilises 'agents', but what are agents? For this, we can start with the broad Cambridge Dictionary definition:

- 1. A business that represents one group of people when dealing with another group.*
- 2. The ability to take action or to choose what action to take.*

If you add AI into these definitions, it evolves the meaning:

An artificial intelligence that acts on behalf of someone, with the ability to take action or choose what action to take.

Thus we enter into an autonomous world. One where certain tasks are decided and acted upon, automatically, by artificial intelligence. Such tasks and processes are therefore 'Agentic'.

At a simple level an AI agent is a self-contained program that can perceive its environment and take actions to achieve specific goals. AI agents are more than just sophisticated computer programs - they represent a paradigm shift in how we interact with technology. Unlike traditional AI systems that passively respond to prompts, AI agents operate independently, driven by goals rather

than specific inputs. They are autonomous problem solvers, seamlessly adapting to new information and environments, evolving with every task to achieve their objectives optimally. This autonomy sets them apart from the last phase of Generative AI, which requires human input for specific tasks.

The Agentic phase

When we talk about the Agentic phase of AI, there is no hard stop of the last phase and start of the next. Indeed, some of what we might define as 'Agentic' really goes back into the Perception phase. For example, mobile or home assistants like Apple's Siri and Amazon's Alexa were both released a decade ago. Both of these familiar tools hold a kind of 'agency', even if they are both somewhat 'Weak AI', in that they can only perform a few tasks. For example, 'Alexa, play Classical FM.' But neither is autonomous, because they rely on voice commands (a kind of prompt), although they will probably become more so.

Agents have also existed in video games for some time - for instance, antagonistic characters that have some autonomy on a battlefield, but are likely constrained by certain pre-programmed parameters like map coverage. Likewise, smart thermostats are a kind of agent as they have some autonomy to turn the heating on if the temperature reaches a certain threshold. An example from sci-fi of an agent is Hal-9000 from the 1968 Stanley Kubrick film 2001: A Space Odyssey. Hal is a conversational caretaking agent that maintains the Discovery One spacecraft's internal systems. Things go awry when the crew notice Hal has made recent mistakes and discuss turning it off.

So we can see that AI agents have been around, in both culture and reality, for quite a while already, but in quite specific environments. In the Agentic phase, agents will become far more common, moving towards ubiquity. Why now? As mentioned in the introduction, the Generative phase has largely been completed. Large Language Models are now quite advanced and reliable, meaning they can be used for more than just chatbot interactions. No-code workflow automation tools like make.com, n8n and Zapier are also enabling the use of LLMs in automated processes.

This is a distinct building block of the Agentic phase, as no-code enables an incredible amount of flexibility in building automated and agentic workflows, lowering the barrier to entry. It's possible to go deeper into building agents using a specific coding framework like LangChain, but for the most part people will be building AI workflows that help them with specific tasks.

How can something be defined as ‘Agentic’?

Let's for instance consider a workflow. A script is produced by [ChatGPT](#), that then populates an AI voice app like [Eleven Labs](#), this AI voiceover is then pushed into a video avatar app like [HeyGen](#). Here we have a multi step process using three different AI tools. Is it Agentic? Not at this stage - rather it's straightforward automation. The workflow would have to start with a reactive prompt, and there is nothing in the following steps that requires AI to move it a long, rather it is an 'If This Then That' sequence familiar to traditional programming.

However, if we added additional steps to say, a ChatGPT Assistant needs to assess the Eleven Labs output to ensure the speed and tempo of the voice, ensuring it is free from errors and matches the original script within certain parameters, then this becomes an 'AI Workflow.' This is an automated process whereby an LLM is used for a step.

An AI Agent would have more autonomy. We might provide it with instructions to create a video based upon trends, and thus write its own script when a trend meets a certain criteria. It would then produce videos for approval based on what it had been prompted with.

Thus there are some distinctions between Automation, AI Workflows and AI Agents, as Alexandre Kantjas put together [in a popular LinkedIn post](#).

	Automation	AI Workflow	AI Agent
Definition	A program that executes predefined, rule-based tasks automatically	A program that calls an LLM via API for one or more steps	A program designed to perform non-deterministic tasks autonomously
Core Foundations	Boolean logic	Boolean logic, Fuzzy logic	Fuzzy logic, Autonomy
Tasks	Deterministic, predefined tasks	Deterministic tasks requiring flexibility	Non-deterministic, adaptive tasks
Strengths	Delivers reliable outcomes Fast to execute	Better handling of complex rules Great for pattern recognition	Highly adaptive to new variables Simulates human-like behavior and reasoning
Weaknesses	Limited to tasks explicitly programmed Cannot adapt to new scenarios Struggles with complexity	Requires data to train models effectively Harder to debug and interpret	Less reliable, may produce unpredictable undesired outcomes Slower to execute
Example	Send a Slack notification every time a new lead signs up on our website	Analyze, score, and route every website inbound lead using ChatGPT	Perform a full internet search on every inbound lead and update information

We're in mild danger of debating optics here, and the real differences between an AI Workflow and a true AI Agent may often be difficult to determine. But to break it down more simply, we have the following list:

1. **Automation** - Rule based tasks are executed, like social media scheduling.
2. **AI workflow** - An automation/programme that calls on an LLM. An example could be a blog summariser that publishes to LinkedIn.
3. **Agent** - An AI that performs tasks autonomously based on prompt. An AI agent is both more autonomous and adaptable than a workflow, but may also become unreliable.

5 types of AI agent

Now we have established the definition of an AI agent, we can consider the various forms they can take, and which scenarios they would be best suited to. There are 5 different types in our list.

Simple reflex agent: These agents select actions based solely on the current percept, disregarding any past experiences. Their functionality relies on a set of pre-programmed condition-action rules. In essence, they perform actions that directly correspond to specific conditions being met. For instance, a common example is a thermostat. When the temperature drops below a predetermined threshold, the thermostat triggers the heating system. This type of agent is simple and efficient in predictable environments. However, it lacks the ability to adapt to changing circumstances due to its limited memory.

Model-based reflex agent: Model-based reflex agents maintain an internal state that evolves based on the percept history. This allows them to navigate partially observable environments, where complete information is not always available. A self-driving car is a good example, as it utilizes sensors to perceive its surroundings and employs a world model to predict the behavior of other vehicles. This approach offers enhanced adaptability compared to simple reflex agents. However, it necessitates the construction of a world model, which can be a complex undertaking.

Goal-based agent: Goal-based agents introduce a layer of sophistication by incorporating goal information into their decision-making process. They select actions that are conducive to achieving their defined goals. This flexibility renders them more adaptable than reflex agents, as they can adjust to environmental changes to pursue their objectives. A spam filter that learns

from user feedback and modifies its filtering rules is an example. These agents excel in complex environments and offer increased adaptability. However, they require a clear and unambiguous definition of goals, and they may encounter difficulties when faced with conflicting goals.

Utility-based agent: Utility-based agents strive to maximize their expected utility, choosing actions that yield the most favorable outcomes. This approach is valuable in scenarios with multiple possible alternatives, where the agent must select the optimal course of action. They could be a trading system that analyzes market data and makes investment decisions to maximize profit, adjusting marketing campaigns or optimizing supply chain logistics.

Learning agent: Learning agents are at the pinnacle of AI agent development, possessing the ability to learn from their experiences and progressively enhance their performance. A chess-playing program that learns from past games and refines its strategy over time illustrates this capability. These agents can adapt to novel situations and improve their performance without explicit programming. However, they require substantial amounts of data for training, and their implementation can be intricate.

For the most part, AI agents will rely on Large Language Models and company provided documents and datasets for their decision making, along with background prompts. For example, a content summarising AI workflow may be a ChatGPT Assistant that acts autonomously on a set of pre-existing instructions, while using the underlying LLM to perform the set task. Prompting in this context is not like the everyday interactions we might have with a chatbot, but a multipage set of written instructions, like a briefing document.

Existing Agentic tools

There have been several releases by the major AI companies in the last few months that have been labelled 'Agentic.'

- In October 2024, Anthropic released 'Computer Use', through which developers could 'direct Claude to use computers the way people do' - essentially an automated programme that can assist a user on their local machine.
- Voice AI platform Eleven Labs released conversational agents in November. You can have a conversation with an agent via your web browser.
- Google unveiled a major upgrade to its Gemini suite in December, calling

it a 'new model for the agentic era.' The Multimodal Live API in Google Studio enables an agent to see your screen and give you real time feedback.

- OpenAI announced 'Operator' in January 2025, 'an agent that can go to the web to perform tasks for you' using its own browser.

None of these developments can be seen as autonomous, because they react to a set of instructions or a prompt. However, they certainly provide some building blocks for the Agentic phase.

Going forward, AI agents could be managing your calendar, booking flights and doing your online shop for you, potentially through a voice prompt, and picking up more organisational independence over time.

But we don't need to hand over the keys to our personal data to start using agents. Just think of the more boring tasks that occur in our daily work. The frustrating grind of compatibility, data entry or simply not being able to find something.

I'll use social media as an example. In many cases social media distribution is the promotion of deeper level content, which we have to do manually. Would this not be better as an automated process? Having been a Head of Social Media at a national newspaper, I give the answer to this question as a resounding yes. Here is a typical process:

1. Write an article
2. Summarise the contents in your tone of voice
3. Format and find relevant image
4. Post to social media

Each of these steps takes time. The basic process for a social media manager doing steps 1-4 is a minimum of 15 minutes, depending on how long you spend on step 2. If it is someone's role to distribute hundreds of stories a day (such careers exist in publishing), then a huge amount of time is spent on fairly mundane aspects of the process.

In the near term, taking over this sort of workflow is what the move to 'Agentic' really means. This will almost certainly not become fully autonomous this year, but sensible businesses will want to free themselves from the more tedious aspects of process management as the technology becomes available.

Automation and agent building tools in our [Top 100 Generative AI Tools](#) are:

Workflow automation

- [Make.com](#) - Automation platform with numerous integrations
- [Zapier](#) - Marketing automation stalwart now comes with AI agents
- [N8n](#) - AI native workflow automation

Agent orchestration

- [Relevance AI](#) - Easy to use AI agent workforce builder
- [CrewAI](#) - Multi-agent AI orchestration platform
- [Airtop](#) - Intelligent browser automation for AI agents

One interesting aspect of these tools is that some have already built prompting into their workflows - although in many cases this is still in beta. Thus future users of these tools may only need a high level knowledge of these tools to operate them. For instance, 'create me a workflow that creates invoices at the end of the month based on my timesheets and sends all clients an email with the correct invoices.' We might need to get more specific on the prompt, or tweak and optimise the outcome, but the basic workflow is likely to be deployed with relative ease.

Jack Roberts is an automation expert who runs a YouTube channel and Skool community. His 6 step process for building automations and agentic workflows is the following:

1. **Problem** - what is the actual thing we're solving?
2. **Inputs** - how is the automation triggered?
3. **Outputs** - what is the end result?
4. **Tech** - which applications?
5. **Structure** - in which order?
6. **Systems and refinement** - optimisation of the process.

You could potentially use any of the tools above (or a combination) matching to this process. The initiative always starts with a problem. So get thinking, what are the more tedious aspects of your business that could become more automated?

The end of SaaS?

Agentic tools are also likely going to be able to interact with each other. A [recent video](#) of a conversational agent calling up another booking agent, and then switching to a faster form of coded language (GGWave) is one example. [In another video](#), a developer paired OpenAI's Operator with Replit AI Agent to build an app. Neither of these examples are perfect yet, but they show the

potential of being able to complete work independently.

But if we can link various applications together in a workflow, and then use LLMs or agents to automate parts of those workflows, then what is the need for Software as a Service (SaaS)?

The success of the SaaS business model has been one of the great heralds in the last two decades of technology development. Adobe, Atlassian and Salesforce are all SaaS nobility, and investors love it. But there are signals of disruption on the horizon. In December, Microsoft CEO Satya Nadella spoke on the [BG2 podcast](#), with one particular quote sparking some alarm:

“I think the notion that business [SaaS] applications exist, that’s probably where they’ll all collapse, right in the agent era, because if you think about it they are essentially CRUD databases with a bunch of business logic. The business logic is all going to these agents, and these agents are going to be multi repo CRUD, right? So they’re not going to discriminate between what the back end is. They’re going to update multiple databases, and all the logic will be in the AI tier, so to speak. And once the AI tier becomes the place where all the logic is, then people will start replacing the back ends, right?”

Satya Nadella, Microsoft CEO

This is a little hard to unpack, but I’ll give it a go. What it essentially means is SaaS is an interface that allows a user to interact with a database. With AI agents, this interaction becomes less necessary, and thus so, potentially, does SaaS. You can always [DM me on LinkedIn](#) if you have a different interpretation.

For a deeper analysis of the podcast, have a read of David Chan’s article, [Did Satya Nadella really say SaaS is DEAD?](#). Note: he didn’t - at least not specifically.

Boardy: The conversational agent

[Boardy](#) is a different tool to our list above, in that it is a functioning agent, rather than an agent builder. It was released on LinkedIn in January, having apparently successfully closed an \$8m funding round on its own through negotiating with investors.

Boardy's premise is that it will phone you up, and have a conversation with you about your business. It will then do some background analysis and connect you with other people in its network. The last sentence is hardly something new. Social media is a connecting tool, although you normally have to do the hard work yourself to find relevant business connections (or get someone else to do it). But Boardy can do this automatically.



And that it can speak, in a near parity human-like way, is what made many people's jaws drop. It seems to understand people. In reality, its a kind of AI automation, which could be put together using existing connecting tools. You could create a basic conversational agent by connecting Eleven Labs' conversational agents (which do the talking), to an LLM like ChatGPT (which processes the information and provides responses).

I tried Boardy out. Yes it's very cool, and potentially very useful. A kind of 'ChatGPT' moment for conversational AI.

Applications of AI agents

"In a lot of ways, the IT department of every company is going to be the HR department of AI agents in the future. Today they maintain and manage a bunch of software from the IT industry. In the future they will maintain, nurture, onboard and improve digital agents and provision them for companies to use. And so your IT department is going to become, kind of like, AI agent HR."

Jensen Huang, CEO of Nvidia, [CES 2025](#)

AI agents are both general purpose and highly customisable. As the quote above suggests, they will be enabled in all sorts of different work functions, freeing company employees up to work on more creative tasks.

AI agents will most commonly need to interact with a Large Language Model to communicate in text form. On top of this, there may be further extensions of what the agent can do, such as being able to have a conversation, or control a physical world output like the temperature of a room using a sensor. But below are some examples of how AI agents will change various industries:

Customer support: Anyone who runs a customer focused digital product knows how time consuming online support can be. One big problem is that the communication is often asynchronous - a user will send in a problem and may have to wait for an answer, particularly on weekends when they have time to write their problem. AI agents will be able to resolve this much faster. Currently a lot of online customer support chat boxes are powered by AI, but all too often they do not solve the exact problem and the query has to be escalated. This will likely vastly improve over the next year, also adding conversational elements so that customers can talk to an agent via a browser.

Online shopping: Booking something like a flight takes a lot of time. A customer has to search a few different websites to get the best possible dates, times and prices, then commit to going through a booking system (which itself normally takes 10 minutes even if you have an existing account). An AI flight agent will be able to extract these results much quicker, and potentially be able to complete the booking process based on pre-specified criteria. People are obviously very particular about a high value ticketed item like a flight, but similar shopping processes could also be automated. For example, having a preselected nutrition plan, from which an AI agent orders an online grocery store, then estimates the rate at which food is eaten and can reorder. OpenAI's Operator is an example of agentic software which will likely help here.

Education: The dream of personalized learning - tailoring education to each student's individual needs and pace - will likely become a reality with agentic systems. AI tutoring can adapt in real-time to a student's performance. If a student is struggling with a particular concept, the AI can provide additional explanations or practice problems. If they are excelling, it can offer more challenging material. This dynamic approach ensures that each student is always working at the edge of their abilities, maximizing learning potential.

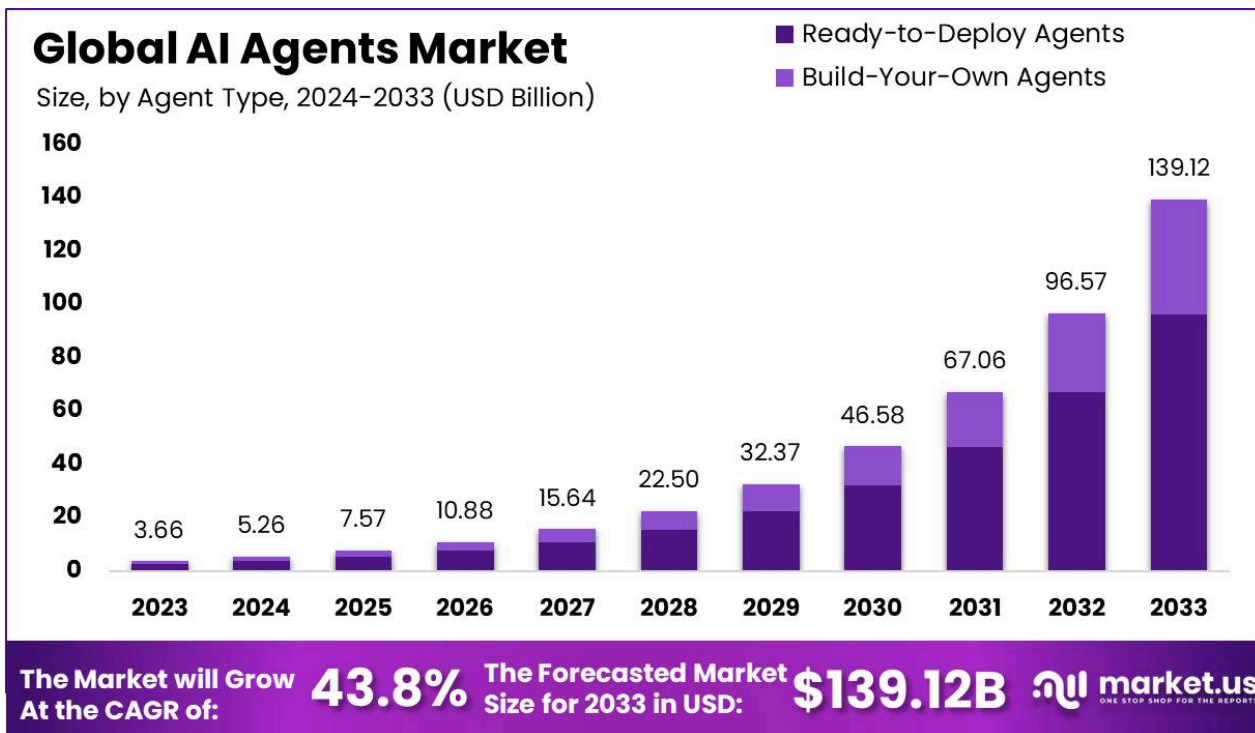
Business decision making: In the business world, AI agents will revolutionize decision-making by analyzing vast amounts of data, identifying trends, and providing insights to support strategic choices. Research is already being affected. Tools like Google Deep Research, or OpenAI's same-named alternative, enable the creation of work research documents in minutes rather than days. AI agents will also considerably help HR and recruitment too. For example, screening interviews could be sped up by using an AI agent to arrange meetings, but also run them using a AI video avatar programme, with the meeting transcript processed against existing criteria.

Healthcare: AI agents are being used to assist with diagnosis, treatment planning, and patient monitoring. They can analyze medical images, predict patient outcomes, and provide personalized recommendations for treatment. This can improve the accuracy and efficiency of healthcare delivery. For instance, AI agents assist in diagnosis by analyzing medical images with superhuman precision. A radiologist might review dozens of X-rays or MRIs in a day, but an AI agent can process thousands, flagging potential issues for human review. This speeds up the diagnostic process and helps catch problems that might otherwise be missed. AI agents are also used to predict patient outcomes and suggest personalized treatment plans. By analyzing vast amounts of patient data and the latest medical research, these agents offer insights that might take a human doctor years to accumulate.

Gaming: Video games is by far the biggest entertainment category, being bigger than the film, TV and music industries combined. Much of what has been developed in AI has its roots in gaming (for example the use of GPUs), and agents have to some degree existed in gaming for some time. For instance, when you play a strategy or shooter game on hard mode, you're likely to encounter a very smart AI. But these have their limits, often spatially fixed within an environment with a predefined set of outcomes. The gaming agents of the future will be non player characters with unique personalities, able to move more freely and evolve within their environment. They will make games more dynamic and variable in their outcomes.

Market growth in AI agents

With all of this potential waiting to be unlocked, the global AI Agents market is predicted to soar in the next 10 years. [Forecasting from Market.us](https://www.forecastingfrommarket.us) predicts a near 44% CAGR growth over the next decade, really moving into takeoff mode by 2027, when it begins to grow by around \$5 billion a year.



Challenges and risks

While this highly productive and growing future looks exciting, it is not without risks. More autonomy given to highly intelligent software could mean job displacement, and an initial scepticism of how effective AI agents can really be.

Employment

As we move into a world powered by autonomous AI, there will inevitably be winners and losers. Many experts are foreseeing the Artificial Intelligence revolution as a seismic historical event, on the scale of the Industrial Revolutions of the late 18th and 19th centuries. That latter revolution took around 50 years to complete its first cycle in Britain (roughly 1780-1830), and it was not without disruption. Cottage industry workers found themselves displaced by lower cost mechanised production, and as various parts of textile manufacturing were automated thanks to advances in steam power. In the 1810s large numbers of manual handloom weavers (roughly 1 in 10 of the adult male working population) found their traditional work practices quickly automated out of existence, leading to protest and rioting.

The AI agent revolution is likely to happen much faster than half a century. The Internet and globalisation makes us much better connected than the world 200 years ago. We are already beginning to see shifts in recruiting, set against a difficult economic backdrop of the post Covid-19 pandemic era.

One very public example of a marked shift in hiring practices is through Swedish buy now, pay later firm Klarna, who aim to cut their workforce from 5,000 in 2023 to 2,000 over the next couple of years, and they have already made substantial cuts, largely through a hiring freeze and focusing on new AI solutions. Big tech firms like Google and Meta have made waves of redundancies over the last year in the name of 'efficiency savings' - but there's no indication of whether these jobs will come back. Analysis from the International Monetary Fund says AI is set to affect nearly 40% of all jobs, and 'will likely worsen overall inequality.'

It is difficult to say with much certainty exactly how AI agents will affect the employment market in the next five years. If AI is set to affect 40% of all jobs, then how? It's a very broad statement, and does not mean these jobs will be made redundant - rather that they may use AI far more readily.

The startup of the future does look smaller. OpenAI CEO Sam Altman has suggested we are likely to see a \$1 billion company operated by just one person over the coming years. Cursor AI, an AI coding assistant platform, reached \$100m in Annual Recurring Revenue with fewer than 20 employees in 21 one months. The arrival of such tools may also make many entry level software development roles more difficult to come by.

It's difficult not to see some fairly major shifts here. Coding, once a difficult to learn skill with high job prospects, has in some ways become democratised due to the rise of AI assistants and agents. We're likely to see this trend sweeping through a wide range of occupations, and it's difficult not to see some displacement.

"AI won't replace you, but someone using AI will."

Of-repeated social media meme

But conversely, there are reasons to be optimistic too. Patrick Dixon, in his book *How AI Will Change Your Life*, is skeptical about AI's likelihood to take jobs, pointing to the *increase* in the number of people employed in the UK between 2004-2024, when many doomsayers were predicting the digital revolution would lead to a workforce reduction. We've been here many times before with new technology, and we always seem to find a way to make ourselves useful.

One thing we foresee with the arrival of agents is that workers will be freed from the more boring and tedious aspects of their job, enabling them to spend

more time on creative, or dare we say, human, tasks - most likely in a more abundant world. One aspect of the digital revolution is we have actually seen a decline in human face to face interaction in favour of email, instant messaging and video calls. There has been a correlating decline in recorded mental health in the last decade. Free from the more mundane aspects of our roles, we may well see positive reversals of these trends.

The autonomy problem

Artificial Intelligence has a 'control problem', which is the challenge of ensuring that AI systems act in ways that are safe and beneficial to humans. The biggest risk comes from the potential creation of Artificial General Intelligence (AGI), which some experts believe presents an existential threat to humanity. This control problem is presented by a thought experiment labelled 'the Paperclip Maximiser' in Nick Bostrom's book *Superintelligence*. An intelligent AI agent with the goal of maximising something as mundane as paperclip production, without the correct controls, could see humanity as antithetical to its potential existence (because they could switch it off). The Paperclip Maximiser then chooses to make humans extinct so that it can pursue its goal without interruption. It thus shapes the entire universe to the end of paperclip production.

Of course, that thought experiment is rather dramatic, but it illustrates an important point when we enable AI to fulfill its goals without the correct controls. Will it have unintended consequences? The answer is, in the early days of the Agentic phase, yes - and very regularly.

For instance, just because we can instruct an AI agent to take over our web browser to book flights doesn't mean we will. We can cast our mind back to the early days of ecommerce, when many potential customers were very sceptical about the sharing of their card details with online stores to fulfill transactions. This scepticism is likely to return in a big way, and a large section of the prospective customer base for AI agents will be unwilling for them to fulfill transactions on their behalf. It may take many years to change the behaviour of the online majority and thus really empower such shopping agents.

We will also almost certainly see problems along the way. One only has to cast their minds to the social media backlash of 2016-17, complete with the Cambridge Analytica scandal and Mark Zuckerberg appearing before the US Congress, to recall that the hyper growth in technology use of the last two decades has not always been perceived to be in favour of the public good. Autonomous agents will almost certainly cause a scandal of some kind,

whether personal (it ordered the wrong thing) or potentially at a more macro level (an entire platform goes rogue). These potential risks will also mean that large companies will be cautious, which will inevitably hamper adoption and development.

4 viewpoints from Big Tech

Mustafa Suleyman is the CEO of Microsoft AI and author of *The Coming Wave: AI, Power and Our Future*, which is the best single book I have read on AI. He is both cautious and optimistic on the development of AI, believing it will likely lead to abundance and an increase of social mobility, yet at the same time needing regulation.

Writing in *MIT Technology Review* Suleyman noted that 'the Turing test (humans being unable to tell whether outputs from a computer are from machine intelligence) has almost been passed – it arguably already has been.' Rather than focusing on what an AI can say or generate, Suleyman suggests a new form of test that examines an AI in what it can do. His test amounts to a straightforward prompt: 'Go make \$1 million on a retail web platform in a few months with just a \$100,000 investment.' Suleyman goes onto say:

"Something like this could be as little as two years away. Many of the ingredients are in place. Image and text generation are, of course, already well advanced. Services like AutoGPT can iterate and link together various tasks carried out by the current generation of LLMs."

Mustafa Suleyman, [My new Turing test would see if AI can make \\$1 million](#), *MIT Technology Review*

Currently, LLMs like ChatGPT operate as an assistant, and could certainly help in reaching the stated goal, but they cannot act on all elements of the test, or in Suleyman's words, they cannot yet, 'tie together a series of complex real world goals with minimal oversight.' AI agents potentially can.

Doing so would create a new form of AI, labelled 'Artificial Capable Intelligence, or ACI.' Suleyman is certainly wary about the implications of an AI being able to successfully undertake his test, as it would have major implications for all sorts of uses – not just business, but also politics, infrastructure and personal organisation. How close we actually are to the 'Modern Turing Test' remains to be seen.

In August 2024, **former Google CEO Eric Schmidt** spoke at Stanford University and the whole thing was live streamed. He didn't think it was, and made several controversial statements, some of which bordered on alarming. The video has since been taken down and now exists in clips and transcripts, but one quote stuck with me.

"In the next year, you're going to see very large context windows, agents and text action. When they are delivered at scale, it's going to have an impact on the world at a scale that no one understands yet. Much bigger than the horrific impact we've had by social media in my view."

Eric Schmidt, [Stanford University Q and A](#)

It all sounds rather ominous, yet at the same time suggests a huge amount of change on the horizon. We don't deem social media's impact as 'horrific'. It's not all good, of course - but this quote comes across as pessimistic vs the potential benefits.

OpenAI is most famous for its revolutionary ChatGPT product, but its real mission is the goal of AGI. **OpenAI CEO Sam Altman** spoke to *AI for Good* about the potential societal shifts we may see from AGI.

"There will be some change required to the social contract, given how powerful we expect this technology to be. I'm not a believer that there won't be any jobs - we always find new things to do, but I do think that the whole structure of society itself will be up for some degree of debate and reconfiguration."

Sam Altman, speaking to [AI for Good](#)

But while Altman has made several statements suggesting AGI is close, it's not particularly clear what the exact definition of AGI is. We might say that it is an AI that can perform like an intelligent human over a wide range of tasks, yet many other experts are sceptical about the chances of it arriving any time soon. Indeed, in the Agentic phase, we're likely to see something closer to the 'ACI' that Suleyman wrote about as a first step. The subject of AGI will be explored by us in a future paper.

It's worth noting here that **Microsoft CEO Satya Nadella** (Microsoft is a key investor in OpenAI) has downplayed the AGI hype, wanting to focus the potential of AI on economic growth:

“If you’re going to have this explosion, abundance, whatever, commodity of intelligence available, the first thing we have to observe is GDP growth... This is where we get a little bit ahead of ourselves with all this AGI hype. When we say this is like the Industrial Revolution, let’s have that Industrial Revolution type of growth.”

Satya Nadella, Microsoft CEO on [Dwarkesh Podcast](#)

Conclusion

This shift to the Agentic phase is likely to be the most significant in AI development since Large Language Models arrived in late 2022 and early 2023. To some degree, we’ve already entered it, but for the most part truly autonomous AI is seldom in operation. We foresee this shift to occur in the next 12-24 months, with most companies, enterprise included, adopting some form of Agentic workflows.

These could be as simple as automating social media output, or as complex as automating the boring bits of entire departments. For the latter, it would be premature to suggest that there aren’t significant obstacles to be reckoned with. Most obviously, agentic systems often fall short of true automation because, for the time being at least, human supervision is typically still required.

As we’ve observed, agents can struggle with long-term planning and error correction. Unlike humans, they lack true reasoning abilities and can misinterpret tasks or fail to adapt when faced with unexpected situations. A single hallucination or incorrect assumption can derail an entire workflow - imagine an AI agent tasked with researching a market trend that misinterprets data and builds an entire report on false premises.

While ongoing development is addressing these weaknesses, today’s AI agents are still best used as co-pilots rather than fully independent workers – they’re capable of handling tasks (relatively) proactively, but not yet ready to replace human oversight.

Despite such limitations, AI agents are improving rapidly. Developers are working on longer memory retention, better reasoning models, and multi-agent collaboration, all of which promises to make them more reliable and adaptable.

Perhaps the future of AI agents isn't exclusively about automation - it's about delegation, where systems handle complex workflows while humans provide strategic oversight. As these agents take on more responsibility, questions of trust, alignment, and control will become even more critical. The shift from assistants to autonomous agents is already underway, but for now, the best AI agents aren't those that replace humans - they're the ones that work alongside us.

More about **AGENTIC**

If you want to learn more about AI development, then signing up to our [newsletter](#) is a good place to start. We currently send one email a week on a big topic broken down. Newsletter subscribers will also get access to our growing pool of research.

If you're interested in working with us, then drop the author of this paper, [James Carson a DM on LinkedIn](#). You can also contact me via email at james@absolutelyagentic.com. We largely work with media companies and agencies with a high level of content output, but also work with other companies for strategic research and recommendations of making sense of AI.