



ARTIFICIAL INTELLIGENCE / Agentic AI / Future of Work / Autonomous Systems / Deep Learning

AI Agents Are the New Employees: The Future of Autonomous Workforces

They don't need breaks or ask for raises. Yet today, they are already handling real and complex tasks that used to require human effort and thinking. This shift isn't something coming - its already happening. And every organisation needs to understand it now, so they adapt and benefit from it instead of being left behind.

By **Maham Shafique** | Agentic AI Engineer | Published: August, 2026

Specialisation: Agentic AI Systems, Machine Learning, Deep Learning, Autonomous Agent Architectures

The Monday Morning That Changed Everything

Imagine a Monday morning at a mid-sized software company. The marketing team comes in, grabs their coffee, and checks their dashboards. But a lot of work has already been done overnight. Emails have been drafted based on customer data, the content plan has been updated using last week's result, a few bugs have been reported from user feedback, and a possible risk of losing key clients has been identified - with suggestions on what to do next. Nobody stayed up to do this. It was handled but AI systems. When the marketing manager reviews everything in the morning, most of it is already good - just small edits are needed. In some cases. She doesn't need to change anything at all.

This isn't something from a sci-fi story - its already happening. Companies of all sizes, from small startups to large global businesses, are using AI systems across different industries like finance, healthcare, and media. This question is no longer whether AI will become part of workforce - it already is. The real question is whether developers, leaders, entrepreneurs, and students - understand it well enough to use it properly. Because very soon, the difference between those who truly understand AI and those who only know the term will become a major gap in careers and opportunities.

As someone who builds AI systems and studies how they work, I want to go beyond a basic explanation of this change. I want to explain what these systems really are, how they work, what they can do well, and where they still fall short. I also want to explore what this shift means for people who will work alongside these systems. For Pakistan's tech community, this conversation is especially important right now. We have a chance to shape how this technology is used here - but only if we take the time to understand it properly and at a deeper level.

"The gap between those who understand agentic AI and those who merely know the buzzword is about to become one of the most consequential divides in professional life."



What an AI Agent Actually Is — Under the Hood

The word 'agent' is being used so much in tech right now that it's starting to lose its real meaning. So let's make it clear. An AI agent, in the technical sense, is a self- operating system that understand its environment through inputs, reasons about what to do using an internal model, takes actions through tools or APIs, observes the results of those actions, and iterates - all aiming for specific outcomes, without requiring a human to supervise each individual step. That last part is the one that matters. The without requiring a human to supervise each individual step is precisely what separates an agent from every AI tool that came before it.

At a basic level, most modern AI agents are built on large language models - the same technology behind tools like ChatGPT - but they are enhanced with three key features that turn them from simple chat systems into independent workers. The first is tool use: the ability to call external functions, APIs, databases, web search, code execution environments, and other systems. The second is memory: both short-term context memory within a task and, increasingly, long-term persistent memory across sessions. The third is a planning loop - the ability to decompose a complex goal into sub-tasks, run tasks one after another or together, evaluate the results, and revise the plan when something does not go as expected.

What's remarkable is that these systems don't need to be taught planning directly - they develop it naturally - it was not explicitly programmed but arose from training on enough human-generated data to learn how to think towards a goal. The deep learning models at the core of modern agents have effectively learned, through exposure to large amounts of text about how people plan things, reason, solve problems, and adapt, how to replicate that process. They are not truly deep thinking about ideas - a debate that continues with full of efforts in both ML research and philosophy of mind - but functionally, for a very wide range of practical tasks, the result look just like real reasoning. And at the level of real-world work, what matters is whether it works the same way.




Figure 1: A conceptual visualisation of an autonomous AI workforce — a central intelligence hub orchestrating specialised AI agents, each assigned distinct roles and tasks, connected through a shared command and data network. Credit: Maham Shafique / PakTech Today



The Architecture of an Autonomous Workforce: How Multi-Agent Systems Actually Work

Even a single AI is quite capable. A coordinated system of AI agents working together is something a different kind of thing - and it is the multi-agent architecture that makes this shift to an autonomous workforce is a major change, not just an interesting data. Understanding how these systems are structured is essential for anyone who wants to build with them or manage businesses that use them.

The Orchestrator. In most well-designed multi-agent systems, there is a main agent that doesn't do the work itself. Instead, it breaks big goals into smaller tasks and assigns them to specialized agents. The Workflow managers (orchestrators) maintain the plan, monitors progress, handles failures from sub-agents, and bring together final outputs. Think of this as the AI equivalent of a project manager or a senior engineer running a sprint - its value is in coordination and judgment, not execution. Frameworks like Lang-graph, Auto-Gen, and Crew AI all provide helper system for building system coordinator AI rank systems, though the design decisions about how to structure the organized levels remain a human engineering problem very much.

The Specialist Agents. These are the workers. Each specialist agent is usually assigned a specific job, a specific set of tools it can access, and a set domain it is expected to focus on. A research agent might have access to web search, academic databases, and ability to condense. A code agent might have access to a Python interpreter, version control APIs, and testing frameworks. A customer communications agent might have access to CRM data, email APIs, and a set of brand guidelines. The specific skill area is not just a conceptual choice - it produces better results than giving a single generalist agent access to everything, because focused context and limiting the tools an AI can use helps reduce mistakes and incorrect outputs.

The Memory Layer. This is the part of system with multiple AI agents that people often overlook and that I find most impressive in terms of engineering. Without memory, agents are stateless - each task starts from scratch, with no awareness of prior context. With a properly designed memory layer - combining databases that help search by meaning, structured knowledge stores for fact - based memory, and tracking past tasks in order - agents begin to behave in ways that feel genuinely employee like. They remember the preferences of specific clients. They learn from past mistakes. They avoid redoing the same work in every session. The memory layer is what transforms an agent from a powerful tool into something that resembles a colleague who learns on the job.

The Feedback and Evaluation Loop. The most solid agent - based systems test themselves automatically of their own outputs - either through a dedicated assessor agent, through automated checks, or through human-in-the-loop review at defined checkpoints. This is analogous to a QA process in a software development workflow, and it is what separates real-world AI system vs demo versions. In my own work building agentic pipelines, the evaluation loop is where most of the engineering time goes - because the quality of feedback decides whether the system improves over time or simply produces consistent mistakes at scale.

"The memory layer is what transforms an agent from a powerful tool into something that resembles a colleague who learns on the job. That shift is not cosmetic. It is architectural."

What AI Agents Can Do Today — Honestly Assessed

A common issue in discussions on this topic is the swinging between the hype and doubt, with little balanced thinking of what agentic systems do well today, where they genuinely struggle, and how those skills are developing



. Let me try to give you that honest picture, because understanding the real capability boundary is the only way to decide the best way to use these systems.

It's clear that AI agents today work well with organized tasks with clear rules and defined work. Software development tasks - code generation, testing, documentation, cleaning up code, code review - are areas where agents are not just useful assistants but real performance enhancers. Tools like GitHub Copilot Workspace and Devin show that AI systems can take general requirements and turn them into working, tested code across different files. Benchmarks like SWE - bench, which test AI on the real issues from GitHub projects, found that top models were able to solve 40% of problems on their own in 2024 - compared to almost none just two years earlier. That growth is very fast.

Data analysis and research synthesis are equally strong areas. An agent with access to web search, a Python environment, and a well-designed prompt can study the market, merge findings from multiple sources, run data analysis on datasets, and produce reports that would previously have taken a junior analyst several days. Customer service and support workflows - responding to critical requests, escalating complex cases, drafting responses, updating CRM records - are already being handled by agentic systems at major companies with clear improvements in how fast issues are resolved and how happy customers are. In healthcare, agents are assisting with clinical documentation, pre-approval workflows, and diagnostic support in ways that free up physician time for the work that genuinely requires human judgment.

Where agent is important just as important to be honest about struggles. Long-horizon tasks with unclear goals and little feedback remain genuinely difficult - agents tend to drift, fits the local context, but globally wrong decisions, and compound errors over extended series of steps. Novel situations that fall outside their training distribution expose the difference between pattern matching and real understanding that is hard to ignore. Tasks requiring real-world physical interaction, deep understanding of people and emotions, ethical judgment in unclear situations, and creative originality - not remixing, but real originality - remain domains where humans maintain a significant and meaningful advantage. These are not small categories. They are, in many ways, the heart of what makes work meaningful. But they represent a less and less of total work time spent in the average knowledge-work company.

The Roles AI Agents Are Already Playing in Organisations

The autonomous workforce is not just a single group of AI systems doing random tasks. Instead, it is made up of different types of AI roles, each designed for specific kinds of work that they handle well. Understanding these roles - and how they change the structure of a company - is what separates people who truly preparing for this shift from those who are only talking about it or going to evens.

The Research Agent. These systems constantly watch incoming data - competitors moves, market developments, academic publications, regulatory changes, news - and synthesise findings into briefings that human decision-makers can act on. A well-designed research agent does not just fetch data; it understands what is relevant to a specific organisation's context, removes irrelevant information, identifies connections between disparate sources, and surfaces insights that might otherwise be missed. Intelligence analysts, market researchers, and competitive strategists are the human roles most directly in dialogue with this capability.

The Engineering Agent. Software engineering agents are perhaps the most mature category of agentic worker today, and they are the ones I work with most directly. These systems can read codebases, understand architecture, write and test code, identify bugs, suggest improvements, and maintain documentation. At the individual developer level, they function as a highly capable pair programmer. At the team level, they are starting to work like independent contributors, handling full features from start to finish. The implication for software



teams - including Pakistan's large population of software freelancers - is both a near-term a big boost in productivity and a long- term concern about entry - level coding jobs.

The Operations Agent. Process automation has always been a technology priority, but these systems add something new: they can handle unusual cases that once needed human judgment. Traditional RPA - robotic process automation - was brittle because it relied on rigid rule sets that broke when the input varied. Agentic operations systems, backed by the flexible reasoning of LLMs, can handle variability in a way that makes them applicable to a far wider range of business processes. Tasks like handling invoices, managing purchasing processes, HR work, compliance checks, and coordinating supply chains are now being done by AI systems at a large scale.

The Creative Production Agent. Marketing, AI systems are changing how content, dressing, and communication work is done. They can handle the whole process - from planning and ideas, to writing, editing, formatting, and sharing the final output. Humans still guide the creative direction at the start and approve the result at the end. The human role is shifting from producer to art director: setting the vision, establishing the standards, and evaluating the output. For some people, this looks like a loss of creative work. But for many creative professionals who used to spend a lot of time on production tasks instead of actual thinking and ideas, it feels more like freedom. The question is whether the company they work for will recognize and capture that value or simply use it as a reason to reduce headcount.

“The human role is shifting from producer to art director. That is not the end of creative work. It is potentially the beginning of more of it — if organisations choose to see it that way.”

The Deep Learning Foundations That Make This Possible

For readers with a technical background, it is worth spending a moment on the deep learning advances that have made the current generation of agentic systems possible - because understanding the basics helps you see both the current limits and how things will likely improve. The AI - driven workforce is not a product of any single breakthrough, it is the coming together of several design improvements that happened to mature at roughly the same time.

The transformer architecture, introduced by Vaswani et al. in 2017 and then expended to create GPT, Claude, Gemini, and their siblings, provided the foundation for language models able to understand complex context over long inputs. The scaling laws that Kaplan et al. documented in 2020 - showing that model performance gets better in a consistent way with increases in parameters, data, and compute - provided the roadmap for building models of enough ability. Instructions tuning and reinforcement learning from human feedback helped researchers align large models with human intent, turning basic language models into systems that can reliably follow complex instructions.

More recently, the development of robust function calling and tool use capabilities - the ability for language models to call external APIs in an organized and dependable way - was the technology makes it possible that turned language models from conversational assistants into genuine agents. Step-by-Step reasoning methods and newer versions of them showed that model's reason is significantly better when they are encouraged to think through problems step by step before producing a final answer - a finding with direct implications for how agentic planning loops are designed. And the emergence of multi-modal models, capable of processing images, audio, and structured data alongside text, has greatly increased the kinds of real - world tasks that agents can handle.

The latest work on AI agents includes better ways to handle long contexts, more reliable tool use with fewer incorrect outputs, improved self-correction capabilities, more efficient fine-tuning methods for domain



specialization, and the development of formal verification approaches for agent safety. Each of these areas, if it matures as the field expects it to over the next two to three years, will expand what autonomous AI workers can reliably do. The progress doesn't follow a straight line but it is steady.

Key Findings: The Numbers Behind the Autonomous Workforce

We are still figuring out, in real time, how AI agents are changing work - but what we are seeing so far is striking. McKinsey's 2024 global survey on AI adoption found that 65% of company reported using generative AI in at least one business function, up from 33% just twelve months prior. Among early adopters of agentic - as opposed to merely generative - AI tools, the productivity gains reported in knowledge work functions averaged 40%, with software development teams reporting the highest figures. Goldman Sachs's research division estimated in late 2023 that AI could automate tasks equivalent to 300 million full-time jobs globally - not eliminate those roles entirely but handle enough tasks to shift how much and what kind of labour is needed.

For Pakistan, where tech jobs in foreign income and drive the local economy, these numbers are real, not just theory. Pakistan has approximately 1.5 to 2 million technology professionals, with software services exports exceeding \$2.6 billion in 2023. The nature of the work that Pakistan's tech exports are built on - software development, data management, customer support, back-office processing - sits almost entirely within the categories where agentic AI is already clearly capable. That's not a reason to panic, but a reason to act fast and plan smartly. The Pakistani technologists who understand how to build, orchestrate, deploy, and maintain agentic systems will not be taken over by those systems. They will be the people others pay to build them.

On the enterprise adoption side, a 2024 study by Gartner predicted that by 2028, at least 15% of day-to-day work decisions will be made on its own by AI agents, up from effectively zero in 2024. Salesforce reported that their Agentforce platform, launched in late 2024, processed over one trillion agent actions in its first quarter of availability. These aren't pilot program results - they are early production results and they're growing very fast.

"Pakistani technologists who understand how to build and orchestrate agentic systems will not be displaced by those systems. They will be the people others pay to build them."

The Human Question: What Happens to the Workforce?

I want to address the question that is in everyone's mind when they read an article like this, because avoiding it would be a form of dishonest thinking. Will AI agents take jobs? Yes, some of them. Certain categories of routine, structured knowledge work - data entry, basic report generation, standard customer support, templated content production, routine code maintenance - will be mostly automated by agentic systems over the next five to ten years. The people doing those jobs today need to understand this and plan accordingly, not because I want to alarm them but as the worst outcome is being caught unprepared.

However, major technological changes in history haven't just removed jobs. These changes usually involve transformation along with net job creation. However the transition can be painful, and the negative effects often hit hardest those workers who already had the least support from institutions. The invention of the spreadsheet eliminated most accounting jobs that existed in 1980. It also created an entire new category of financial analyst, business intelligence professional, and data scientist that did not exist before. The internet eliminated travel agents, video store clerks, and classified ad salespeople. It created search engine optimisers, social media managers, cloud infrastructure engineers, and UX researchers. Overall, employment increased, but the change required adjustment.



The agentic AI transition will follow this pattern, but there's one key difference: its being adopted faster and affecting more types of work than ever before. The shift of agriculture took hundreds of years. The industrial revolution over a long period of time. The digital revolution over about 10 years. The agentic AI revolution is unfolding over years, possibly less. The institutions - educational systems, labour markets, rules and regulations, training systems for professionals - that have historically helped workforces adapt to technological change were not built for transitions at this speed. That is perhaps the most structural challenge right now, and it is one that governments, universities, and employers all need to take ownership of rather than waiting for someone else to solve.

Safety, Oversight, and the Engineering of Trust

Any honest discussion about the AI - driven workforce must include safety and oversight issues, and as someone who builds these systems, I feel a strong responsibility to explain them clearly. Deploying AI agents in important work processes - ones where mistakes have real costs, whether financial, reputational, legal, or human - should not be decided based on excitement instead of serious thinking about technology. The engineering of trust in agentic systems is itself a serious discipline, and it is one that the field is still developing.

The key safety principle in designing AI agent systems is keeping their impact as soon as possible: agents should request only the permissions they need, prefer actions that can be undone, and ask a human when you are confident enough about the right choice. This sounds simple in principle. In practice, designing systems that reliably adhere to it across complex, dynamic real-world environments is genuinely hard. It requires careful prompt engineering, robust evaluation frameworks, sandboxed execution environments for high-risk tool calls, comprehensive audit logging, and human-in-the-loop checkpoints at the right granularity - not so often that they reduce autonomy, and not so rarely that mistakes build up before being caught.

The second safety area is alignment: making sure agents follow the goals people truly intend, not similar - looking ones that can fail in edge cases. This is not a solved problem in the research literature, and it is not a solved problem in production systems either. It is an active area of engineering work that every team building serious agentic systems is grappling with. The company that will deploy agentic workers most successfully and most responsibly will be the ones that invest in alignment and oversight infrastructure as seriously as they invest in capability and scale. Skipping safety to move faster may seem reasonable in the short term but can look disastrous in hindsight.

What This Means for Students and Early-Career Professionals Right Now

I am a computer science student myself, and I want to speak directly to the students and the early-career professionals reading this, because the advice you are getting from people who formed their career intuitions in the pre-agentic era is . The value of your education has not decreased - if anything, it has increased, because the foundations of computer science, machine learning, and software engineering are exactly the foundations you need to work at the frontier of agentic AI. But the application of that education needs to be different from what it was five years ago.

Do not think of yourself as someone who writes code. Think of yourself as someone who architects systems - systems that increasingly include AI agents as first-class components alongside human contributors. Build things with agent frameworks now, before you graduate. Not tutorials, not demos - real systems that solve real problems, even small ones, so that you develop genuine intuition for where agents work well and where they break down. Learn to read and understand ML research papers, because the gap between the research frontier



and production systems in this field is narrower than in most areas of technology, and the people who can translate between the two are extraordinarily valuable.

Develop taste and judgment about when to use agents and when not to. One of the most common and costly mistakes in agentic system design is deploying agents in situations that would be better handled by simpler, more deterministic approaches - because the engineering team was excited about the technology rather than rigorous about the problem. The ability to say “an agent is not the right tool here” is just as valuable as the ability to build one. And invest deeply in the skills that remain distinctively human - systems thinking, ethical reasoning, stakeholder communication, creative judgment, leadership - because these are the capabilities that will matter most in a world where the routine cognitive tasks are handled by machines.

What's Next: The Horizon of the Autonomous Workforce

Over the next three to five years, if we look at things realistically, autonomous systems are going to become much more common in everyday work. They won't just handle simple tasks anymore - they will be used in more areas, do more kinds of work and slowly start taking on responsibilities that matter more. In short, they will become a bigger and more trusted part of how work gets done. The next generation of agentic systems will be multi-modal by default - seeing, hearing, and acting on visual and audio inputs alongside text. They will have much better at thinking through long term problems step by step, enabling them to manage complex, months-long projects rather than just hours-long tasks. They will be more tightly integrated with the the physical world using robots and embodied AI platforms, allowing automation to extend from digital work into the real, physical parts of organizations.

Most importantly, we will likely see long lasting AI agents that keep working inside organizations, building up knowledge over time and becoming more skilled in their specific areas. The difference between a human expert who has been at a company for ten years and a junior hire who joined last month is not primarily a difference in raw intelligence - it is a difference in accumulated context, pattern recognition, and relationship knowledge. Agentic systems with persistent memory and continuous operation are beginning to develop analogous advantages, and this will make them qualitatively more valuable over time, not just incrementally better.

For Pakistan's technology ecosystem, the window of opportunity is real, and it is open right now. The company, researchers, and individual practitioners who develop genuine depth in agentic AI - not just familiarity, but deep, hands-on, engineering-level depth - will be positioned to lead rather than follow this transition. The AI agents are already here, already working, already reshaping what company look like and what human work means within them. The only question that remains is who builds the future and who inherits it. For everyone reading this who is willing to do the work to understand these systems at the level they deserve, the answer to that question is still entirely open.

About the Author

Maham Shafique is an Agentic AI Engineer and Machine Learning and Deep Learning expert currently pursuing her degree in Computer Science. She works at the forefront of autonomous agent architectures, building production-grade agentic systems across a range of industries. Her technical focus spans large language model orchestration, multi-agent system design, deep learning model development, and the safety and evaluation frameworks that make autonomous AI deployable at scale. She is passionate about equipping the next generation of Pakistani technologists with the deep technical literacy needed to lead — not follow — the autonomous AI transition.

Keywords: Agentic AI; Autonomous Agents; Future of Work; Multi-Agent Systems; Large Language Models; Deep Learning; Machine Learning; AI Workforce; LangChain; AutoGen; Pakistan Tech; Autonomous Workforces



References

- [1] Vaswani, A. et al. (2017) *Attention Is All You Need*, NeurIPS 2017. arXiv:1706.03762
- [2] McKinsey Global Institute (2024) *The State of AI in 2024: Generative AI's Breakout Year*, McKinsey & Company. www.mckinsey.com
- [3] Goldman Sachs (2023) *The Potentially Large Effects of Artificial Intelligence on Economic Growth*, Goldman Sachs Global Investment Research.
- [4] Anthropic (2024) *Building Effective Agents*, Anthropic Research Blog. www.anthropic.com/research
- [5] Gartner (2024) *Top Strategic Technology Trends for 2025: Agentic AI*, Gartner Research. www.gartner.com