# Artificial Intelligence – Are Expectations Being Met?

When OpenAI released ChatGPT in November 2022, it sparked the AI revolution we're living through today.* For the first time, cutting-edge artificial intelligence wasn't locked away in research labs or corporate data centres—it was a free website anyone could use. The numbers tell the story: ChatGPT reached one million users in five days and 100 million monthly users within two months, outpacing the growth of Netflix, Instagram, and TikTok combined.

## What Made ChatGPT Go Viral

ChatGPT succeeded because it delivered two things at once: genuine practical value and a captivating story about machine intelligence that felt almost superhuman.

The practical side was obvious. People could draft emails, debug code, explain quantum physics to a teenager, or brainstorm business ideas—all through a simple chat box that required no training to use. But the emotional hook ran deeper.

The interface felt conversational and human. You typed a question, and it responded with fluent, thoughtful answers that seemed to come from a single, coherent mind rather than a statistical algorithm. News outlets amplified the spectacle: ChatGPT passing the bar exam, writing poetry, composing essays indistinguishable from student work. Headlines proclaimed "AI that can do anything," and millions believed it.

This tapped into something ancient in the human psyche. We've spent millennia asking fundamental questions: Why do we exist? What makes us conscious? Is there something greater than ourselves? Religion offered divine answers. Science pursued empirical explanations. Philosophy debated the nature of mind and soul. Now AI has become the latest vessel for these existential longings—a digital oracle that might unlock cosmic truths or even surpass human intelligence. At the same time, building such systems satisfies a god-like impulse: the desire to create beings superior to ourselves.

Science fiction has been rehearsing these fantasies for decades. Stanley Kubrick's HAL 9000 in 2001: A Space Odyssey (1968) showed us eerily intelligent machines. Ex Machina (2014) explored an AI so convincing it manipulated its human observers. The Matrix (1999) questioned whether reality itself might be simulated. Steven Spielberg's A.I. Artificial Intelligence (2001) gave us a robot child yearning to be loved. These stories primed us to expect AI would change everything about being human.

## The Reality Behind the Hype

AI is genuinely transformative. It's accelerating drug discovery, automating complex business processes, and reshaping industries from healthcare to finance. But the hype has outrun the reality. Predictions of mass unemployment and machines that "understand" the world like

humans don't match what these systems actually do. Surveys suggest 99% of people fundamentally misunderstand AI's capabilities and limitations.

To assess whether expectations are being met, we need to look past the romance and examine what AI actually is—not thinking machines, but extraordinarily sophisticated pattern-matching systems.

## What ChatGPT Actually Does

ChatGPT is a conversational interface built by OpenAI that connects users to large language models, or LLMs—specifically, versions of GPT (Generative Pre-trained Transformer). Similar tools from Google (Gemini), Anthropic (Claude), Meta (Llama-based assistants), and Elon Musk's xAI (Grok) work on the same principle: massive neural networks trained to generate human-like text in real time.

## The Engine Under the Hood

At its core, ChatGPT runs on transformer-based neural networks with billions or trillions of parameters—essentially adjustable weights that determine how the system processes language. These models are pre-trained on vast datasets: books, websites, code repositories, scientific papers, and more. The training process is unsupervised, meaning the system learns by trying to predict the next word in a sentence across trillions of examples.

After this initial training, engineers fine-tune the model using Reinforcement Learning from Human Feedback (RLHF). Human reviewers rate thousands of responses, teaching the system to prioritise answers that are helpful, accurate, and safe. When you type a prompt, it's sent to clusters of thousands of graphics processing units (GPUs) that generate a response word by word—or more precisely, token by token—based on probabilistic patterns learned from the training data.

## How It Creates Answers

Traditional search engines retrieve and rank existing content. ChatGPT does something fundamentally different: it synthesises new text from scratch.

For example:

- *Prompt*: "Explain quantum computing in simple terms."

- *Response*: "Quantum computers use quantum bits, or qubits, which can exist in multiple states simultaneously, unlike traditional bits that are either 0 or 1..."

The system hasn't "looked up" this explanation. It has reconstructed it from billions of similar explanations it encountered during training, blending patterns to produce something that reads as original and coherent.

This generative approach makes ChatGPT excellent at writing, brainstorming, summarising, and explaining. But it also leads to a well-documented problem: "hallucinations," where the system confidently states falsehoods or invents details that sound plausible but are wrong. It can also struggle with mathematics or logic without specialised plugins to support it.

Hallucinations are compounded by the "loss of context memory". When a chat uses up the context memory either through large files or lengthy chats, it no longer just invents details but loses the oversight and creates contextless, flawlessly worded gibberish. The whole chat becomes useless.

## The Intelligence Illusion

The term "Artificial Intelligence" suggests these systems think the way humans do. They don't. A more accurate name would be "Pattern Synthesiser," but that doesn't inspire the same awe—or attract the same venture capital.

## Fluency Versus Understanding

LLMs learn to predict the next token in a sequence based on statistical patterns in enormous datasets. They "stitch together sequences of linguistic forms according to probabilistic information about how they combine," as researchers put it, without any direct access to meaning or the physical world.

The result is text that is syntactically flawless, logically structured, and emotionally attuned—qualities we associate with intelligence. Studies comparing ChatGPT conversations to human dialogue have found that the AI scores higher than people on measures of analytical thinking, cognitive clarity, and social language. The prose is unusually polished and coherent, which reinforces the impression of a sharp, attentive mind.

But this is form, not substance. The system has learned what smart explanations sound like without necessarily grasping the underlying concepts. It's "mathematically induced fluency," as

critics describe it—outputs that look insightful because they mimic the patterns of expert discourse, not because they reflect deep understanding.

## Why People Are Fooled

Most people don't know that ChatGPT is essentially a next-word prediction engine scaled to industrial size. Many assume it retrieves pre-written answers from a database or follows fixed scripts. This misunderstanding reinforces the myth of an all-knowing oracle rather than a probabilistic text generator.

RLHF makes this worse in a sense, because it explicitly optimises for style: confident, polite, well-structured answers that match what we culturally expect from a knowledgeable expert. The system has been trained to sound intelligent, which is not the same as being intelligent.

## The Problem Extends to Audio and Video

The same illusion of authenticity that makes ChatGPT sound intelligent is now appearing in generative audio and video. Tools like ElevenLabs can clone voices with startling accuracy. Platforms like Runway, Sora, and others produce video that looks real. As these technologies improve, the line between reality and fabrication blurs—and that creates a systemic problem.

## The Erosion of "Seeing Is Believing"

Deepfake technology—using generative adversarial networks (GANs), diffusion models, and voice synthesis—can now produce audio and video that fools experts, especially after compression or re-posting degrades the original. Real-world harms are already documented: fraudsters using cloned voices or deepfake video calls to trick employees into transferring millions of pounds. Detection tools struggle to keep pace as generation quality improves.

Research shows that repeated exposure to deepfakes reduces public trust in media and institutions, even when individual fakes don't fully convince viewers. People begin to doubt everything.

## A Crisis of Truth

The deeper issue isn't just that fiction looks real. It's that the fundamental assumption underpinning modern society—"seeing is believing"—stops working. We risk entering a post-

truth environment where anything can be dismissed as fake, which is just as dangerous as believing every fabrication.

Analysts treat hyper-realistic synthetic media—text, audio, video—as one of the central challenges of generative AI. It supercharges misinformation, fraud, and reputational attacks while undermining the shared evidentiary basis (photographs, recordings, eyewitness accounts) that democracies and legal systems depend on.

Just as polished language makes LLMs sound intelligent, vivid synthetic media makes fakes look and feel real. We're moving into a world where distinguishing reality from fabrication is both cognitively and technically harder—and that's one of the main structural risks of generative AI.

## The Verdict: Expectations Versus Reality

AI is delivering transformative tools that boost productivity and unlock new creative possibilities. But it isn't delivering the existential breakthroughs or job apocalypse that headlines promised. The fervour stems from our projections—ancient desires for oracles, gods, and answers to life's biggest questions—not from what the technology actually does.

The challenge now is recalibrating expectations. AI is a powerful instrument, not a panacea or a threat to human existence. By understanding its real capabilities and limits, we can harness its benefits without succumbing to hype-driven disillusionment or the erosion of trust that deepfakes and misinformation bring. The AI boom will endure, but only if we see it clearly.