

# 1

# UNVEILING GENERATIVE AI: A NEW FRONTIER

Okay, let's go back to basics. What is generative AI (artificial intelligence)? How does it work? And what is the technology capable of? You'll find out in this chapter as we delve under the hood of GenAI, tug at a few wires, and examine what the heck's going on in there.

I really want this chapter to give you an inspiring feel for the many things that GenAI is capable of – but also to drive home the point that GenAI isn't just about ChatGPT (Chat Generative Pre-trained Transformer). Sure, ChatGPT is a prime example of GenAI (and it certainly hoovers up the majority of GenAI's press), but there's a lot more to GenAI than ChatGPT . . . as you'll find out in this chapter.

## What Is Generative AI? A Quick Explanation

Actually, let's start by defining artificial intelligence (AI) in its broadest sense. The term "AI" refers to computer algorithms that can effectively simulate human cognitive processes, like learning, decision-making, and problem-solving.

GenAI is a groundbreaking subset of AI – the cutting edge of the cutting edge – that is able to create new content based on patterns and structures it has learned from existing data. Like any AI, GenAI tools are given enormous amounts of data to learn from (what's known as "training data"). They learn

from the training data, and then use the patterns or rules that they've learned to create new content that's similar to, but not exactly the same as, the data they have been trained on.

### **An example or two**

A good example is DALL-E 2 – the text-to-art platform that allows anyone to generate artworks. Or, of course, there's ChatGPT, the language model that can create text based on conversational text prompts. (You'll find a much bigger list of GenAI tools in the Appendix, by the way.) These systems learn from huge training datasets – ChatGPT, for example, was trained on vast amounts of text from the internet, including web pages, articles, and books.

Text and images are perhaps two of the best-known uses of GenAI so far, but the technology is capable of so much more. With GenAI you can generate product designs, computer code, music, video, voices, and even entire visual worlds. We'll talk more about GenAI's capabilities later in the chapter, but to whet your appetite, imagine being able to create unique video game worlds rendered in real time, or have a book written *just for you*, or have your favorite celebrity read you today's news. The possibilities are mind-blowing.

The capabilities are already quite impressive, but in the future, GenAI systems will be able to create pretty much anything that humans can. And this, in turn, means GenAI can turn anyone and everyone into an author, musician, computer programmer, filmmaker, or other type of creator.

### **How GenAI differs from the AI we are used to**

AI is increasingly part of the world around us, including the search results you get on your phone, your conversations with Alexa, and the movie recommendations that Netflix serves up on an evening. So what makes GenAI different to these “traditional” AI tools? (I realize it sounds strange to refer to “traditional AI” when it's hardly been around that long, but I do so to distinguish between the AI that we're already used to in everyday life and this new evolution of AI systems. One technical term for traditional AI is “discriminative AI.”)

Traditional AI systems also learn from large amounts of data, but they deliver a different output. Traditional AI systems are used to make *predictions* based on existing data. And we use those predictions to help us make better decisions, at work and in everyday life. This could span anything from listening to new music on Spotify, and viewing recommended products on Amazon, to identifying which of your company’s customers are most likely to buy a certain product.

This new wave of GenAI tools goes even further, by *creating new content* based on existing data. In other words, GenAI isn’t just about simulating human cognitive processes like decision-making and problem-solving – it’s about simulating *human creativity*.

To further illustrate the difference, imagine you’re playing computer chess. The computer knows all the rules, can predict your moves, and makes its own moves. It’s not inventing new chess moves; rather, it’s selecting the right move based on existing strategies. That’s traditional AI – it’s like a master strategist who can make smart decisions within a set of rules. And it does it very well. But GenAI? Well, that could, in theory, come up with new ways of playing chess that we haven’t invented yet.

So, traditional AI excels at pattern recognition, while GenAI excels at *pattern creation*. Which is very cool indeed.

That said, GenAI and traditional AI aren’t mutually exclusive. They could be used in tandem to provide even more powerful solutions. For example, a traditional AI could analyze user behavior from your company’s website, and then a GenAI could use the analysis to create personalized content for users.

## How Does Generative AI Work?

Think of it like learning to draw by looking at many pictures. After looking at many pictures, you try to draw something new on your own. GenAI does something similar: it “looks” at a lot of data (text, visual, or other), learns the patterns within that data, and then tries to create something new that fits those patterns.

So, in very simple terms, GenAI is like an artist or writer that has studied lots of existing works and then tries to create its own work based on what it has learned. This process is powered by complex algorithms that mimic how our brain works in order to learn from data and identify patterns.

That's the super-simple explanation. Let's get a little more technical.

### **Introducing machine learning and neural networks**

We'll talk more about the evolution of GenAI in Chapter 2, but as a quick primer, GenAI grew out of a field of AI study and practice called "machine learning" – indeed all of the AI we see today is based on machine learning. While traditional computer algorithms are coded by a human to tell a machine exactly how to do a particular job, machine-learning algorithms are able to make decisions based on what they learn from the data. The more data they're fed, the better they get at this process.

Another term you'll need to get to grips with is "neural networks," as this is the core technology that GenAI is built upon. A neural network is essentially an advanced machine-learning model inspired by the way human brains work. While a less complex machine-learning model may need some human intervention in the process, a neural network has the ability to learn and make decisions by itself, and can even learn from its own errors – rather like the way a human learns through a process of trial and error.

Here's how a neural network works:

- **Training:** neural networks learn through a process called "training." During training, the neural network is fed a lot of data (which could be text, images, or whatever), and it learns to identify patterns and relationships in the data.
- **Learning:** as the neural network is exposed to more and more data – and we really are talking about vast amounts of data – it gradually gets better at identifying patterns and understanding the underlying rules that govern the data.

- **Layers:** neural networks are organized into layers, and each layer is responsible for identifying different types of patterns. The initial layers might identify simple patterns, and as we move deeper, the layers are able to identify more complex patterns.
- **Generative models:** GenAI often uses specific types of neural networks called “generative models.” In addition to recognizing patterns, generative models are able to generate new data that is similar to (but not exactly the same as) the data they were trained on.
- **Input and output:** once the generative model is trained, you can give it an input (e.g., a partial image or a text prompt), and it will generate an output (like a completed image or a piece of text) based on what it learned during training.
- **Randomness:** GenAI introduces a certain amount of randomness in the generation process, which means that it can produce slightly different outputs each time, even when given the same input over and over again.

### Examples of generative models

I just mentioned generative models, the neural networks that enable GenAI to create new content. Here are some examples of generative models used in GenAI applications:

- **Large language models (LLMs):** by gobbling up large amounts of text, LLMs learn the semantic relationships between words and use that data to generate more language. An example of an LLM is GPT-4, created by OpenAI, which powers the ChatGPT tool.
- **Generative adversarial networks (GANs):** these work by pitting two competing algorithms against each other, one tasked with generating data that resembles its training data and the other tasked with trying to tell whether the output is real or generated. This type of model is typically used to create images, sounds, or even video (including deepfakes).
- **Variational autoencoders:** this is a type of model that learns how data is constructed. It’s often used to generate synthetic data.

- Diffusion models: these work by adding random data (known as “noise”) to the data it’s learning about, and then figuring out how to remove it while preserving the original data. This way, the model learns what’s important and what can be discarded. Diffusion models are most commonly used in image generation.
- Foundation models: this is something of an umbrella term for models that are pre-trained on a broad range of data at an enormous scale and then fine-tuned for specific tasks. These are large-scale, general-purpose models that are the foundation of LLMs and other GenAI tools.

With different models able to do different things, it’s clear that GenAI has a wide range of applications. Which brings us to. . .

## What Can Generative AI Do?

In Part 2, we’ll delve into how GenAI is already being used in practice (with lots of real-world examples), but for now, I really want to make the point that GenAI is capable of creating all sorts of content.

## GenAI is so much bigger than ChatGPT

Of course, ChatGPT gets a lot of attention – and for good reason: it’s an impressive tool, and creating text is a major application of GenAI. Which is probably why “ChatGPT” and “generative AI” have almost become synonymous terms. But GenAI is capable of so much more than writing.

Take Coca-Cola’s Masterpiece advert, for example. A collaboration between human artists and GenAI, the ad brings many of history’s greatest artworks to life on the screen in a way that’s never been done before. GenAI has also been used to create a new Beatles song by rebuilding partially recorded lyrics by John Lennon, combined with new material by Paul McCartney.

But GenAI isn’t just relevant to creative content. It’s also being used to speed up drug discovery, with one UK company recently announcing that it’s created the world’s first AI-generated immunotherapy cancer treatment.

## UNVEILING GENERATIVE AI: A NEW FRONTIER

Then there's "generative design" – an emerging field where GenAI is used to create designs and production processes for new products. In one instance, General Motors used generative tools created by Autodesk to design a new seat-belt bracket that's 40% lighter and 20% stronger than its existing components.

And we can't talk about GenAI without talking about "deepfakes," which blur the lines between reality and fiction by making it appear as if real people have done or said fake things. GenAI is the technology behind deepfakes, and it's only going to get better at creating fake – but incredibly realistic – content. Deepfake Tom Cruise was one of the earliest and most famous examples (Google it if you haven't seen it – it's uncanny). More insidiously, potential candidates on both sides of the 2024 US presidential elections have starred in deepfakes designed to discredit them.

As if spreading false information wasn't bad enough, there are also outright criminal applications of GenAI, including attempts to extort money by staging hoax kidnappings using cloned voices, and, in one memorable example, fraudulently scamming money by posing as a company CEO.

As the technology has become more sophisticated, GenAI's capabilities have expanded greatly. To summarize, GenAI can be used to generate:

- Text
- Images
- Video
- Music and sound
- Designs
- Data analysis
- Computer code
- Synthetic data
- Research
- Video games and virtual worlds

Let's explore each usage in a little more detail.

## **Text – writing text that’s (pretty much) indistinguishable from a human writer**

Alongside ChatGPT, there are other generative text tools like Google’s Bard and Meta’s Llama. These tools can be used to write anything from essays and articles, to plays, poems, and novels.

Here are some of the incredible things that generative text tools can do for you:

- Writing assistants: helping you draft emails, reports, and other written materials efficiently.
- Creative writing: assisting authors in coming up with unique story ideas, characters, plotlines, or even drafting entire narratives.
- World-building: aiding authors in creating vivid and immersive worlds for their narratives, with tools that help outline histories, geographies, and cultures of fictional settings.
- Writing prompts: offering a variety of writing prompts to help writers overcome creative blocks and generate new ideas.
- Grammar and style corrections: offering real-time suggestions for improving the grammatical accuracy and stylistic fluency of written content.
- Translation assistance: help with translating texts to different languages while retaining the nuances of the original content.
- Accessibility: providing features like speech-to-text to help people with disabilities draft texts.
- Song writing: creating song lyrics, even in the style of your favorite artist.
- Recipe writing: coming up with recipes that suit your parameters (e.g., “a dinner recipe that uses butternut squash, feeds a family of four, and takes no longer than 45 minutes”).
- Plain-English explanations and text summaries: helping you understand a complex topic in plain English, or providing a succinct summary of a long piece of text.
- Multiple-choice questions: generating quiz or test questions for a variety of topics, to beginner, intermediate, or advanced levels.

- Travel itineraries: aiding you with everything from finding flights and accommodations to suggesting popular attractions to visit. Even where to get a good meal, based on your dietary needs and budget.

### **Images – creating images to your specifications**

Many GenAI tools – such as Midjourney or Stable Diffusion – can take a natural language prompt (i.e., regular, conversational language) and use it to generate a picture. Tell it you want an image of a two-headed dog wearing an Elvis costume flying a spaceship into a black hole and watch as it (or something close to it) appears before your eyes.

This can be useful for:

- Art creation: assisting artists in generating unique artworks by understanding and mimicking different art styles.
- Style transfer: enabling artists to apply the styles of famous artworks to their creations.
- Automated sketching: assisting in automatically generating sketches based on the descriptions provided, aiding artists in the initial stages of their work.

### **Video – from editing to creation**

While not yet as advanced as text or image generation, tools are beginning to emerge that allow us to create and edit video simply by describing what we want to see.

Typical uses include:

- Enhanced editing: automating intricate editing tasks including cut detection, color correction, and subtitle generation, thereby streamlining the post-production process.
- Dynamic special effects: allowing real-time rendering of high-quality special effects, motion tracking, and 3D modeling, significantly reducing the time and effort involved in post-production.

## GENERATIVE AI IN PRACTICE

- 3D modeling: facilitating the creation of 3D models and environments for integration into videos, helping to create immersive visual experiences.
- Sound design and synthetic voice: enhancing audio experiences through AI-generated music compositions and synthetic voices. Which brings us to. . .

### **Music and sound – from voices to songs**

GenAI tools can create human-like voices (voice synthesis), allowing computers to speak words. They can also create music and sound effects.

In practice, this can be used for:

- Harmony suggestions: proposing harmonious chord progressions and melody lines.
- Generative music: creating entire musical pieces automatically based on a set of input parameters.
- Mood-based creation: offering tools that help composers create music tailored to specific moods or emotions.
- Voice synthesis: generating realistic voices that can read text naturally, providing a human-like auditory experience.
- Language and accent variability: offering a wide range of language and accent options, helping to create diverse and inclusive auditory experiences.
- Customizable voices: providing tools to create customizable voices, allowing users to personalize their auditory experiences.

### **Graphic design and generative design – streamlining design work**

GenAI offers huge potential for graphic designers and product designers to streamline the design process – from automating aspects of visual design to creating multiple product design variations.

Some of the ways this may prove useful include:

- Design suggestions: aiding designers by providing design suggestions.
- Color palette suggestions: providing suggestions for color palettes that work well.
- Stock image integration: offering seamless integration with stock image libraries, giving designers easy access to a wide range of resources.
- Architectural design: assisting architects in generating building designs optimized for specific criteria like materials cost, energy efficiency, and spatial requirements.
- Product design: helping designers to come up with optimized and innovative product designs by exploring a vast number of design alternatives in a short time.
- Fashion design: assisting fashion designers by proposing new clothing designs based on current trends, historical data, and designer inputs.

### **Data analysis and reporting – democratizing data analysis**

With GenAI tools, you don't need to be a data scientist to interrogate data. From generating automated reports to providing useful predictions, there are lots of ways GenAI can help businesses get the most out of data.

In terms of automated reporting, GenAI offers:

- Real-time reporting: offering tools that generate reports in real-time, providing up-to-date information for decision-making.
- Visual representation: creating visual representations of data, such as graphs and charts, to make reports more understandable and engaging.
- Natural language summaries: generating summaries of reports in easy-to-understand language.

And in terms of predictive analysis, GenAI can deliver:

- Anomaly detection: identifying unusual patterns in data that may indicate potential issues or opportunities.

- Forecasting and customer insights: utilizing historical data to make accurate forecasts about future trends, and identify customer preferences.
- Simulation and modeling: creating detailed simulations based on data analysis, allowing organizations to explore potential scenarios and make better decisions.

### **Coding – computer programming made simpler**

As well as ChatGPT, tools like Microsoft's GitHub Copilot and Amazon's CodeWhisperer make it easy for anyone to generate computer code with very little technical knowledge.

This can span:

- Code generation: creating code based on high-level descriptions or partially completed code.
- Code review: assisting in reviewing codes by identifying bugs, security vulnerabilities, and providing suggestions for optimization.
- Automated testing: creating tests automatically based on the software's specifications, helping to improve the software's reliability.

### **Synthetic data – solving privacy issues and more**

GenAI makes it easy to create entirely synthetic data for use in training other AI models. And for everyday businesses, synthetic data can help them overcome the privacy and data security obligations that come with real-world data.

In practice, this may include:

- Data augmentation: creating additional data through techniques like image augmentation to enhance the training of AI models.
- Privacy preservation: generating synthetic data that preserves the privacy of individuals by not using real personal data.
- Balanced datasets: creating balanced datasets that help in reducing biases in AI models, promoting more fair and equitable AI systems.

- **Diverse data:** generating diverse datasets that help researchers explore a wide range of scenarios and conditions.

### **Research – making discovery easier and faster**

There are many ways GenAI can aid the research process, with a few notable examples being:

- **Drug discovery:** speeding up the drug discovery process by predicting potential drug candidates through the analysis of vast datasets.
- **Simulation environments:** creating realistic simulation environments based on synthetic data, aiding researchers in conducting experiments and studies.
- **Rare event modeling:** assisting in modeling rare events through synthetic data, helping researchers explore scenarios that are difficult to study with real data.

### **Video games and virtual world generation – making immersive content more responsive**

From virtual reality (VR) environments and immersive video game worlds to the wider concept of the metaverse, we're increasingly interacting with immersive, highly realistic content. Designing such content is a complex task that can be accelerated with the help of GenAI. Here's how:

- **Procedural generation and design:** creating intricate environments, realistic terrains, and adaptive storylines, fostering unique and immersive gameplay experiences.
- **Character design and animation:** generating natural character animations, and crafting consistent character personalities to enhance storytelling and immersion.
- **Personalized user experiences:** tailoring games and immersive environments to individual preferences and behaviors, fostering a more personalized and engaging user journey.
- **Automated content moderation and security:** monitoring and maintaining the safety of immersive environments, using automated systems to detect and manage inappropriate content.

## The Impact of Generative AI

GenAI is clearly capable of amazing things and has the potential to transform business and everyday life. I'll explore examples from specific sectors across Part 2, but for a broader overview, turn to Chapter 3.

The rapid evolution of GenAI is also forcing us to confront some difficult questions, and fast. For example, what happens when we reach a point where it's impossible to tell the difference between what's real content and what is generated by AI? There are no easy answers right now, but we'll do our best to delve into such thorny issues in Chapter 4.

And then there's the question of how all this will affect human jobs, as machines take on increasingly creative work. That's the subject of Chapter 5.

## Key Takeaways

In this chapter, we've learned:

- GenAI is a groundbreaking type of AI that is able to create new content based on patterns it has learned from existing data.
- While traditional AI simulates human cognitive processes like decision-making, GenAI simulates *human creativity*.
- GenAI is made possible by neural networks – an advanced machine-learning model inspired by the way human brains work.
- ChatGPT may have made GenAI famous, but GenAI is much bigger than ChatGPT. GenAI can be used to create music, sound, and video, write computer code, design products, create immersive worlds, automate data analysis and reporting, and more.

Before we move on to the impact of GenAI and the issues surrounding it, let's briefly explore the evolution of GenAI, from the early days of computing to now.