IN THIS CHAPTER

- Selecting the right GenAI model
- Understanding GenAI capabilities
- Creating with GenAI
- Grasping its many uses
- Conquering fears and piloting opportunities

Chapter 1

Mapping the Lay of the Generative AI Land

Welcome to the exciting world of Generative AI (GenAI)! This chapter is your starting point in understanding the vast landscape of GenAI and its transformative capabilities. Whether you're a curious beginner or a tech enthusiast, you'll find the information here to be an accessible guide to the basics of GenAI. You can easily build on these skills through practice, regular use of an AI application, or by returning to this book from time to time to enhance your skills further.

So, What Exactly Is Generative AI?

You can think of AI (short for *artificial intelligence*) as incredibly sophisticated software. Although it doesn't behave like any other software ever made, it is still software. Illustrations depicting AI as robots reflect the difficulty in drawing AI software in a way everyone will instantly recognize. But the robot is actually mindless hardware, and the AI is the "smart" brain-mimicking software installed to enable it to function in ways we consider to be intelligent in a non-organic sense.

Technically speaking, GenAI refers to a subset of artificial intelligence technologies that use sophisticated *natural language processing* (NLP), neural networks, and *machine learning* (ML) models to generate unique and humanlike content. It belongs to a classification of AI called *Large Language Models* (LLMs), which analyze huge amounts of data in numerous languages including human languages, computer code, math equations, and images.

LLMs typically have a substantial number of parameters, which are numerical values used to assign weight and define connections between nodes and layers in the neural network architecture. Parameters can be adjusted to change the weights of various values, which in turn, changes what the model prioritizes in the prompt and data and how it interprets various data points, words, and connections.

Imagine you have a recipe for making a cake, and the recipe is your GenAI model. The ingredients — like flour, sugar, eggs, and butter — are like the data points, words, and connections in the model. Now, the amount of each ingredient you use (how many cups of flour, how much sugar, and so on) are like the weights of various values in the GenAI model or GenAI application.

Just as you might adjust the ingredients in your cake to make it sweeter or fluffier by adding more sugar or an extra egg, you can adjust the parameters in a GenAI model to change what it focuses on and how it interprets the information it's given. If you want your GenAI to pay more attention to certain words or data points, you increase their *weight* just like adding more chocolate chips to your cake if you want it to be extra chocolatey. This way, the GenAI model, like your cake, turns out the way you want it to, based on what you prioritize in the recipe.

LLMs use parameters to predict the next word in a sequence — meaning they predict the word most likely to follow the words in your prompt, and then the word that most likely follows its first predicted word, and so on until the model believes it has finished the most probable pattern. It generates images in much the same way by predicting the image that follows your description in the prompt. The models can complete the process incredibly quickly. For example, LLMs like GPT-3 and GPT-4o developed by OpenAI are capable of processing billions of words per second. It is the speed of its response, the appearance of nuanced understanding, and its fluid use of natural language that gives GenAI interactions a humanlike feel.

REMEMBER

However, GenAI and LLMs are not human and do not think — again, they predict. It's a very complicated prediction process, to be sure. Nonetheless, it is a prediction. And if anything happens to tilt its predictive capabilities, nonsense ensues. You can see one example of that in Figure 1-1, which is an OpenAI incident report about an adjustment they made to the model resulting in ChatGPT responding to users in incomprehensible gibberish.

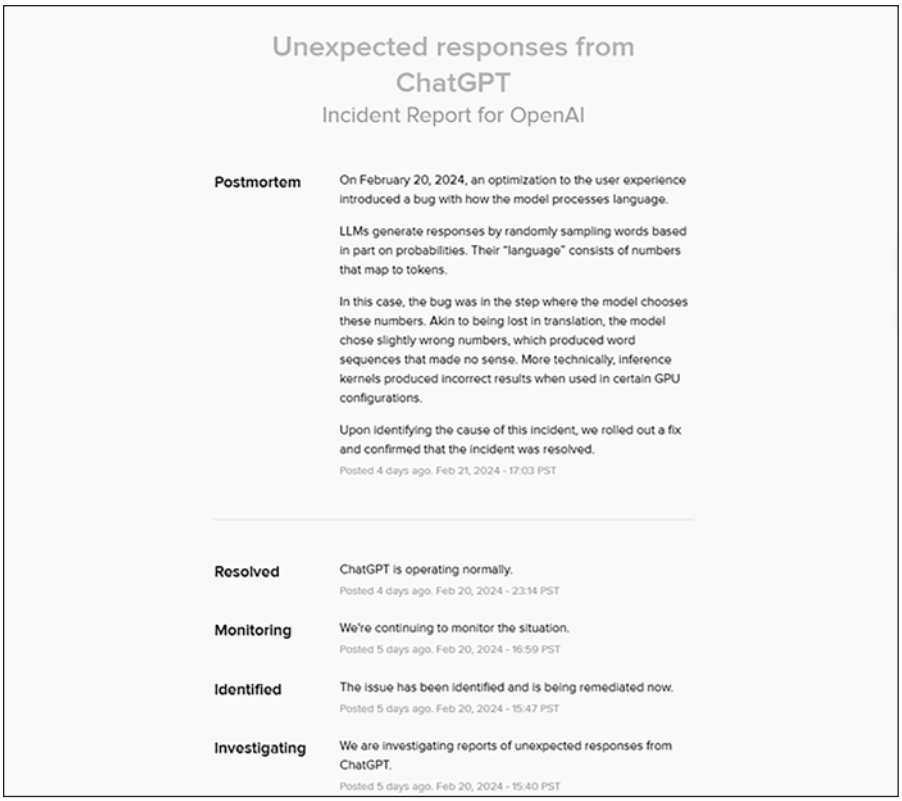

FIGURE 1-1: A routine effort to optimize ChatGPT resulted in its producing gibberish in response to users' prompts.

Source: OpenAI incident report

GenAI VERSUS VIRTUAL ASSISTANTS

AI models and applications are the software driving the robot or the autonomous car or whatever form it's given in the corporeal world. But strictly speaking, AI has a digital form. Because of that, it can be squeezed into almost anything, and many a vendor does exactly that. You'll find various types of AI are embedded or otherwise at use in all sorts of products and services. However, not all AI is the same.

Here are the main differences between GenAI apps like ChatGPT and virtual assistants like Siri, Alexa, and Google Assistant.

(continued)

(continued)

Virtual assistants:

This class of AI runs on a proprietary mix of technologies in a blend developed by their respective corporate owners. Certain components, such as machine learning, deep learning, natural language processing, smart search or search engines, and speech synthesis make the assistants appear and sound much like ChatGPT.

However, their responses are more limited than GenAI models. People typically use these to retrieve answers to common questions or perform uncomplicated tasks like "where is the nearest pharmacy?" or "play a song by Taylor Swift" rather than to generate original answers.

GenAI models (specifically ChatGPT in this comparison):

This class of AI runs on a single AI model, meaning on one version or another of Generative Pre-trained Transformers (GPT) AI models. GenAI is a broad category of AI that includes models capable of varying capabilities such as generating text, images, or computer code or some combination of these.

People typically use GenAI web apps, but some mobile apps and a few wearable devices are available as well. But in all cases, the apps run on a single GenAI model.

Unveiling the BIG Secret to Working Successfully with GenAI

If you remember nothing else I've written in this book, you must remember what I tell you in this section. For here is the big secret — the master key — that you need to make GenAI models work at the level you need them to perform. If you don't grasp this, GenAI will likely appear to you to be nothing more than a fascinating toy or a tool that falls far too short of your expectations.

TIP

In a nutshell, GenAI generates outputs that appear to be original thoughts or images from a computer, rather than results produced by very advanced, contextual predictive software. GenAI retrieves words or images pulled from a database and repurposes them into a new response. The big secret is that the humanlike feel in the "conversation" is an illusion. You are not having a conversation with a machine. It doesn't understand a word you wrote in your prompt.

REMEMBER

Current GenAI models don't think or create things *per se*, but instead *generate* new things from parts of old things found in its database. (The term "things" in this context being images, videos, numbers, or text, depending on the GenAI application you are using.) A GenAI output is the model's best prediction of what you are seeking. In an oversimplified explanation of a complex technology, GenAI seeks to complete a pattern that you began with your *prompt*, which is your question or command as entered into the prompt bar on the GenAI's user interface (UI). In other words, GenAI predicts what letters, words, or images are likely to follow those that are in your prompt. Its predictions are based upon comparison to patterns that exist within its training dataset and/or datasets to which it was subsequently given access.

Think of GenAI outputs as the result of repurposing or remixing information that the model has access to in datasets, including the following:

- Data it is exposed to in its training database along with any additional data provided in subsequent fine-tuning.
- Data added in system messages or prompts.
- Data added via methods such as *retrieval-augmented generation* (RAG), which is a tactic to enhance accuracy, relevancy, and reliability by adding external sources to the GenAI's database.

  RAG combines the strengths of both *information retrieval AI,* which is a set of algorithms that retrieve contextually relevant information from huge datasets, and GenAI, which uses neural networks and machine learning models to generate new content. It might help to think of RAG as GenAI that is augmented by more traditional information retrieval AI, or retrieval AI for short.

WARNING

Since GenAI generates outputs that are the result of its remixing or repurposing of information, it has no concept of true or false, fact or fiction. GenAI can accurately define these terms, but it does *not* understand their meaning. It doesn't understand anything you wrote in the prompt or that it wrote in its response. It only appears to understand terms and concepts. This is an illusion. This is why you must always factcheck its work.

REMEMBER

GenAI responses are limited to the confines of the data it has access to. Put another way, if its training data were a mound of Legos and there were no end caps in that mound, GenAI would build its outputs without end caps. It would not know that end caps exist at all. In the same way, it does not know fact from fiction unless those labels are applied to specific data points in its dataset. But, if a falsehood is labeled as fact, GenAI will unquestionably accept it as fact. It still doesn't understand the difference.

To illustrate this analogy, I wrote a caption first and then used it as a prompt in Azure OpenAI Studio DALL-E playground (Preview). The result is the stunning concept illustration you see here in Figure 1-2.

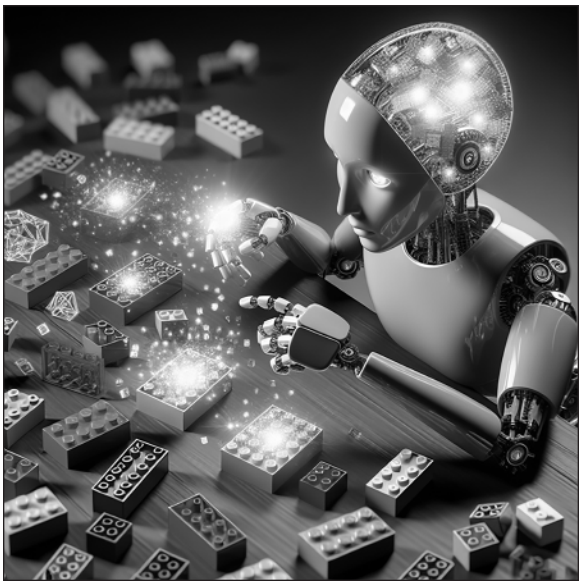

Art generated by Azure AI Studio DALL-E playground.

FIGURE 1-2: If data were Legos, GenAI could only build things with the Lego pieces it has access to, and it is completely unaware that any other types of Lego pieces exist.

GenAI can repurpose and remix only the data it has access to, which is a major reason why GenAI outputs can be highly reliable or totally false or something in-between. The data itself can be insufficient — in one way or another — to provide the foundation or elements for the model to generate an accurate answer. Outdated data from an aging training dataset and data limited to too few perspectives or examples are common issues, but there are many others.

When — not if — outputs are wrong, people call them *hallucinations.* It's unclear why no one calls them lies, falsehoods, or simply errors, but in any case, you cannot assume that GenAI outputs are solid enough to bet your life or business on without doing some serious factchecking first.

REMEMBER

While GenAI does consider context when it analyzes the words in your prompt, it does not understand you or what you said in the prompt. This is why you must not confuse GenAI with General AI, also known as Artificial General Intelligence (AGI). AGI does not yet exist outside of science fiction movies, books, and TV shows. Yet

some people are so in awe of GenAI capabilities that they are sure this must be it — the thing from the movies that's going to take over the world! This is not that.

Understanding the Infamous Finger Problem and Other GenAI Quirks

Perhaps the most wondrous thing about using GenAI is the delicate dance between human and machine that begets something neither would have made alone. But once you move past the first exhilarating moments of viewing GenAI marvels, you'll begin to see a few cracks here and there.

For example, it is common for GenAI models to draw people with six or more fingers on one hand. This is typically because the patterns it sees in its data is of multiple fingers on one human hand. No clear pattern emerges of there being just five fingers on one hand, so GenAI can't predict how many fingers it needs to generate.

Essentially GenAI is parroting the answer from its database. It doesn't understand the question or the answer; therefore, it does not know to draw only five fingers. Instead, it looks for patterns in hands depicted by images or text in the datasets to which it has access. But the pattern of the total number of fingers is unclear. Images in most databases that GenAI models use typically show hands in different positions wherein only some fingers are visible or fingers from two hands or more are intertwined. GenAI cannot therefore see a consistent pattern of the total number of fingers per hand. However, if you were to ask the model how many fingers are on a hand, it will almost always tell you that there are five. Even though it gives you the right answer as to the number of fingers, it does not understand its own reply and, therefore, still doesn't *know* the answer.

Data pattern inconsistency and the resulting probability prediction error is why you can end up with too many or too few fingers in any image GenAI generates. This is often the reason for other issues in images and videos that GenAI creates such as errors in shadowing or movement.

REMEMBER

Although GenAI is impressive, its reasoning is limited. In fact, it's extremely difficult for GenAI to reason at all. To overcome this shortfall and make it more powerful, add one or more humans to the mix and you'll soon see real magic in the result. It is the collaboration between you and this extremely sophisticated software that will take you to the goals you seek.

Figuring Out How to Work with GenAI — It's All About Your Prompts

Here's the thing: Natural human language is a computer language now. In the case of GenAI, this means that the machine still works like a machine and the human like a human, but they can now interact through a computer language that everyday, non-programmer types of people can understand and use.

However, you, the human, still must think like a machine to get the most out of GenAI. Ask any computer programmer how important it is to think like a machine while programming — and this is true regardless of their choice of programming language, be that JavaScript, Java, HTML/CSS, SQL, Python, English, or French.

And why is changing how you think important? Because you are not having a conversation with GenAI. You are giving instructions (and, yes, even when your prompt is a question, it is an instruction) on what you want the model to produce, much like any programmer does. You must think beyond the language to the depths of the result you seek. The value of a programmer is not their computer language knowledge, although that is important too, but the problem-solving ability that they can then convert into language that renders the precise solution the programmer wants to produce. This is how you need to think and work with GenAI models, too.

Your prompts need to be more concise and detailed than the typical conversations you have with another human. For one thing, you cannot make assumptions that a listener will automatically fill in common details because GenAI often doesn't know those details. Despite appearances, GenAI does not think and doesn't truly understand your prompt; many of the natural assumptions you make in speaking to another human will not work in the same way in interactions with these models.

TIP

The GenAI winning formula: Machine speaks like a human. Human thinks like a machine. The better you get at telling GenAI what you want, the better it'll get at giving you what you need. It's all about practice.

Why GenAI appears so human

From crafting sentences to conjuring up images, composing music, or creating synthetic data, GenAI is a master in making something instantly that can often readily pass as human made.

The interesting thing is that its outputs ***are*** human made in some sense. GenAI can be thought of as a creative tool like an artist's paints, crayons, and pencils. Those

items produce images first imagined in the artist's mind and executed by the artist's hand and skill. Similarly, GenAI delivers outputs according to the user's vision and skilled prompting.

Further, much of the data that GenAI models learn from is generated by humans. But it learns far more from this information than you might imagine. It also learns the habits, attitudes, biases, and other human attributes behind the text, audio, and image data that it consumes.

For example, GenAI models have been known to be "lazy" in the summertime around peak vacation periods — meaning GenAI models may produce less content in response to a prompt than usual. GenAI may even tell a user to get the information for themselves. Sometimes a GenAI model also responds slower than normal, announces a delay, or makes excuses.

Such actions aren't due to a bug or a flaw in the system. The AI is merely mimicking human behavior. Models learn human behavioral patterns along with data patterns from their training dataset. They make no distinctions in the values of the information versus the behavior and so are likely to distribute both, or either, in their outputs.

GenAI can also deliberately lie and act angry or sad or cheerful for the same reason. It may even appear to ignore you from time to time. On the flipside, GenAI models tend to perform better when given a virtual reward or a compliment. Again, all of this is just mimicry of the human behaviors it has learned. It's important to remain aware of such idiosyncrasies when using GenAI. Strategically playing into these GenAI quirks can level up the responses you pull from it.

Depending on the model you're using, inputs and outputs can be in text, images, and/or audio forms. Unlike traditional AI, which analyzes, makes decisions, and delivers outputs drawn from data, Generative AI can repurpose information to create seemingly original outputs in a conversational or artistic manner. But it can also plagiarize and pillage the works of other humans. You must always check its outputs for grievous and potentially liable or dangerous behaviors.

WARNING

A key point to remember is that if you use GenAI you are legally liable for what it does. It is not a "separate legal entity . . . responsible for its own actions" as Air Canada once argued trying to defend itself in a court case after ChatGPT gave one of its customers incorrect information. You'll find more information in the discussion on responsibilities in Chapter 3. (And if you're curious about the Air Canada story, you can access it here: `www.bbc.com/travel/article/20240222-air-canada-chatbot-misinformation-what-travellers-should-know`)

Realizing the human influences behind Generative AI's abilities

It's important to distinguish between generating and creating. GenAI "creates" text or images by generating a response from repurposed information based on its prediction of the "best" match to your prompt. GenAI does not create in the truest sense, which dictionary.com defines as "to cause to come into being, as something unique that would not naturally evolve or that is not made by ordinary processes."

GenAI works mostly by making predictions, which technically is an ordinary process in that predictions are a common thing that people and analytics do. However, GenAI's prediction processes are quite extraordinary in that they exist at a level never before achieved. By machine standards this achievement is extraordinary because it generates a new response as opposed to a regurgitated response or picking one of a limited number of "canned" responses. By human standards, GenAI's performance is extraordinary because it can analyze huge amounts of data and respond in a conversational manner or with a newly generated image in seconds or minutes.

TIP

One way to remember the difference is to think "To generate is AI, to create is human or human and AI."

But make no mistake, GenAI is not as humanlike as it appears. Another distinction is in motivation. Humans are motivated to create; some even feel driven by their passions. By contrast, GenAI is not motivated to generate anything. Ever. It doesn't get hungry, thirsty, lonely, inspired, emotional, cold, hot, uncomfortable, dedicated to a cause, politically activated, or otherwise stimulated so there's no reason for it to do anything at all.

You must provide the vision, the passion, and the impetus in a prompt. Then it will try to generate whatever that is for you. Otherwise, it will sit idle for centuries — or however long its supporting hardware and electrical power exists. That's why no one need worry whether GenAI will take over the world.

However, everyone should worry about the humans using GenAI to take over the world. As a tool, GenAI is neither good nor bad. But its users can be either or both. It is the blend of human and AI capabilities that makes GenAI models perform so uniquely and wonderfully. And sometimes comically or poorly.

Discovering the Differences in GenAI Models and Options

GenAI interacts through natural language and generates new content by repurposing data into new outputs. They are most commonly used in areas that require fresh ideas and original output, such as customer service, graphic design, digital media, entertainment, software development, and writing.

However, they can also be used in specialized tasks for a variety of industries such as healthcare, pharmaceuticals, life sciences, manufacturing, and the financial sector.

This section offers a breakdown of specific GenAI models and their corresponding outputs.

Image outputs:

- **DALL-E 2:** This AI model can convert textual descriptions into detailed images or artistic creations, demonstrating the power of language-based image synthesis.
- **StyleGAN 3:** This model is known for generating high-resolution, photorealistic images of subjects such as human faces, animals, and vehicles, offering customization options. It's also used to animate images.
- **Stable Diffusion:** This GenAI model specializes in generating lifelike images, videos, and animations derived from textual descriptions and visual prompts.
- **Imagen:** Trained to understand and interpret image-text pairings, this GenAI system excels in crafting images from textual cues and performing neural style transfers.
- **Adobe Firefly:** A GenAI tool designed for converting written descriptions into visual content, Adobe Firefly aids in the creation of artistic and creative imagery.
- **Midjourney:** This GenAI tool is adept at converting textual prompts into distinctive and captivating artwork very quickly.

Text outputs:

- **ChatGPT:** Developed by OpenAI, this advanced chatbot can generate text that is coherent and indistinguishable from human conversation across various topics.
- **OpenAI Codex**: This model specializes in generating and completing code based on natural language prompts, forming the backbone of tools like GitHub Copilot.

- **HuggingChat**: This is an open-source AI chatbot created by Hugging Face, providing a ChatGPT-like experience using the Open Assistant Conversational AI Model for dialogue-based engagements.

Audio outputs:

- **Jukebox:** Another creation by OpenAI, Jukebox composes music across different genres, illustrating GenAI's capacity to craft musical pieces.
- **PaLM 2**: A Google-developed transformer model that excels in generating multilingual content and performing coding tasks.
- **AudioCraft**: This suite includes MusicGen, AudioGen, and EnCodec, three distinct models that work in tandem to produce authentic audio and music based on textual descriptions, providing an avenue for crafting rich and captivating auditory content.
- **Project Music GenAI Control:** A nascent tool from Adobe Research, this generative AI specializes in music creation and refinement, enabling artists to spawn musical pieces from text inputs and adjust the resulting audio.

Video outputs:

- **Stable Diffusion:** This model employs diffusion techniques to generate photorealistic images, videos, and animations from textual and visual prompts.
- **Neural Radiance Fields (NeRFs):** This novel neural network approach can be used for creating 3D visuals from 2D image data.
- **Synthesia:** An AI video generator tool that transforms textual input into video content, featuring AI-driven avatars and voiceovers for simplified video production.

Multimodal inputs and/or outputs (generates more than just text):

- **Copilot AI:** Made by Microsoft, this model aims to boost workplace efficiency by offering chat-based interfaces for information retrieval, composing emails and summaries, crafting images from textual descriptions, and programming in multiple coding languages.
- **ChatGPT 4o (omni):** This model allows multimodal inputs and generates multimodal outputs. Additionally, the availability of specialized GPTs in the GPT Store can be used to add capabilities. For example, Image Generator can be used within ChatGPT to create images to illustrate its textual output.

- **Gemini:** A suite of generative AI models from Google DeepMind and Google Research, designed with multimodal functionalities to process text, images, audio, video, and programming codes.

REMEMBER

GenAI models are always learning and getting better, and there's a new one popping up all the time. Which one you pick depends on what you need it for, how much you want to spend, and how easy it is for you to use.

Checking Out Practical Uses of GenAI

GenAI models are the subject of many news stories, water-cooler talks, zoom meetings, and online chats. It seems nearly everyone has an opinion on where this technology is leading. Some predict doom and gloom while others expect rainbows and riches. The real story about GenAI is far more practical and realistic than the talk surrounding it.

Following is a list of popular practical uses today:

- **Content creation**: GenAI models like Claude and ChatGPT 4o are being used today to assist authors, scriptwriters, screenwriters, speechwriters, and other creatives in generating stories, speeches, character dialogues in games and movie scripts, marketing collateral, ads, blogs, websites, and even entire books like this one. They can be used to enhance human creativity, or just handle the background research, planning, storyboards and character tracking.
- **Visual arts**: Image generators such as DALL-E, Midjourney, and Stable Diffusion can create photorealistic images or artworks from textual descriptions, aiding artists and designers in visualizing concepts and creating new art forms. Specialized models can create everything from storyboards to short videos and select gaming or movie scenes.
- **Search and knowledge assistance**: Generative AI is being integrated into search engines and virtual assistants, transforming them into more powerful knowledge assistants. When ChatGPT or a similar GenAI application is embedded in a search engine like Bing, you can usually choose between reading the provided helpful narratives of the search results or reviewing a list of sources and related items. Some search engines like Perplexity are built from the ground up on GenAI and even provide a list of sources used to generate search summaries.
- **Customer service**: ChatGPT can interact with customers to provide support, offer solutions, and facilitate service processes like returns, making customer service more personalized and efficient.

- **Education**: GenAI can disrupt traditional education models by enabling customized learning plans and assessments for students, moving away from one-size-fits-all approaches to competency-based learning progression.
- **Media and journalism**: While GenAI cannot perform investigative reporting, it can assist journalists by providing background information, context, and faster dissemination of news stories. It can also manage basic, fact-only reporting such as sports scores and daily stock market analysis.
- **Legal and data analysis**: Lawyers can use GenAI to draft legal documents. Staff can use it to conduct analyses of mounds of case-related documents and evidence to quickly derive insights and timelines and write reports, allowing the attorneys to focus on strategy and deeper insights. However, the final legal documents must be overseen and edited by lawyers or their appointed clerks and paralegals as GenAI does makes mistakes.
- **Marketing and advertising**: Marketers and advertisers can leverage GenAI to produce content and ads rapidly, keeping up with emerging trends, making better market fit, and enabling continuous delivery cycles.
- **Smart automation**: As virtual assistants like Siri and Alexa become integrated with GenAI, they will become more intelligent and versatile, capable of understanding and anticipating user needs. Integrating GenAI with other software and eventually with autonomous AI agents will complete a smart automation cycle from user request to task completion. For example, instead of Google Assistant just giving you a list of restaurants nearby, with GenAI and other software or app integrations, it can also book a reservation or place a to-go order for you too!

Separating Gen AI Fact from Fiction

With any emerging technology, myths and misconceptions can arise. It's important to separate the hype from reality. For instance, while Generative AI can automate certain tasks, it doesn't mean it will replace all human jobs.

One common misconception is that Generative AI will lead to mass unemployment. While it's true that AI can eliminate jobs much like automation tools have done, it also creates new job opportunities and roles that require human oversight and creative input. For example, AI-generated content still needs human curation to ensure quality and relevance.

Another myth is that Generative AI can independently create high-quality content without human intervention. The quality of AI-generated content heavily depends on the input and guidance it receives from human users. GenAI is a tool that amplifies human potential, not a replacement for human creativity.

Table 1-1 lists some of the most common myths about GenAI today and the realities to match.

TABLE 1-1 Common GenAI Myths versus Reality

Myth	Reality
GenAI can take over for human creativity.	GenAI isn't about to steal the job of artists, inventors, innovators, photographers, videographers, content managers, medical researchers, scientists, or other professionals in a wide array of disciplines. It's more like a helpful sidekick that can pitch in with ideas and content, but it still needs a human boss to direct and oversee its work to make sure the final product makes sense and shines. And that the patient survives.
GenAI speaks all languages perfectly.	Think of GenAI as a language student; it's pretty good at languages it's been taught, but it's not a polyglot prodigy. Training it to understand different tongues takes a lot of data and effort, so it's not equally slick in every language.
GenAI is totally fair and neutral.	Just like people, GenAI can pick up biases from the stuff it learns. To keep GenAI fair, humans have to step in and guide it, sort of like teaching it good manners and enforcing laws.
GenAI is a jack-of-all-trades.	Generative AI tools are not uniform but rather tailored for specific functions, with each having unique pros and cons. For example, tools such as ChatGPT3 are optimal for language-related tasks, whereas DALL-E and Midjourney are great for creating images, underscoring the necessity to match the tool to the task. Multimodal models like ChatGPT4 appear to be good at a lot of different types of inputs, but they have their drawbacks too. For instance, ChatGPT 4o isn't great at making graphs and infographics — at least not yet.
GenAI is a wild, untamed beast.	GenAI might be the new kid on the public block, but that doesn't mean it's a wild card. The scary stories are often more fiction than fact. Like any technology, it's all about how you use it.
GenAI will flip business on its head overnight.	GenAI is pretty awesome for whipping up content and making things more personal and efficient, but it's not a magic wand for business. Its superpowers work best when they're tailored to specific tasks and business goals.
GenAI is just a step away from thinking on its own.	Let's get this straight: GenAI is clever at making stuff based on patterns, but it's nowhere close to thinking or feeling like a human. It's smart, but not that kind of smart.
GenAI will take over in the future.	It's a tool, not an overlord. It's also been overhyped and massively adopted too early for user skills to catch up, which means that we're probably in a bubble that will burst soon. Never fear, this is a typical cycle for new technologies, albeit on steroids this time. GenAI is here to stay. Learn the skills now to secure a job now and in the future. It's just going to take a bit for the dust to settle so people and companies can see where this tech works best.

REMEMBER

Generative AI is a rapidly evolving field. Its capabilities are expanding, and with them, our understanding of what is possible. The journey through the landscape of GenAI is just beginning, and the path ahead is filled with opportunities for innovation and growth.