IN THIS CHAPTER

- Accessing Gemini for the first time
- Starting to query Gemini
- Exploring the chatbot interface
- Evaluating subscription plans
- Discovering multimodality

# Chapter 1

# Welcome to Gemini, the Multimodal AI for Everyone



Gemini represents both the chatbot and the underlying engine that powers all other Google AI tools, accompanied by algorithms and multimodal technologies that extend its generative capabilities.

Because this chapter orients those who have not yet used Gemini, Google's main AI tool, I start with the basics — that is, from when you first access the platform. I do this so you can start using Gemini right away.

So, if you have already created a Gemini account and taken your first steps with the chatbot interface, feel free to skim through the chapter to pick up any information you may have missed that could be important for using the platform.

# Jumping into Gemini

WARNING

AI evolves rapidly; while this book strives for accuracy, the screenshots in this chapter may differ for their current interface.

You get started with Gemini by making your first query. Follow these steps:

1. **Connect to Gemini via your web browser (preferably Chrome, but it is not essential) using the address `gemini.google.com`.**

   If you have not yet created a Google account or have not logged into your browser, you'll see a notification message like the one in Figure 1-1.

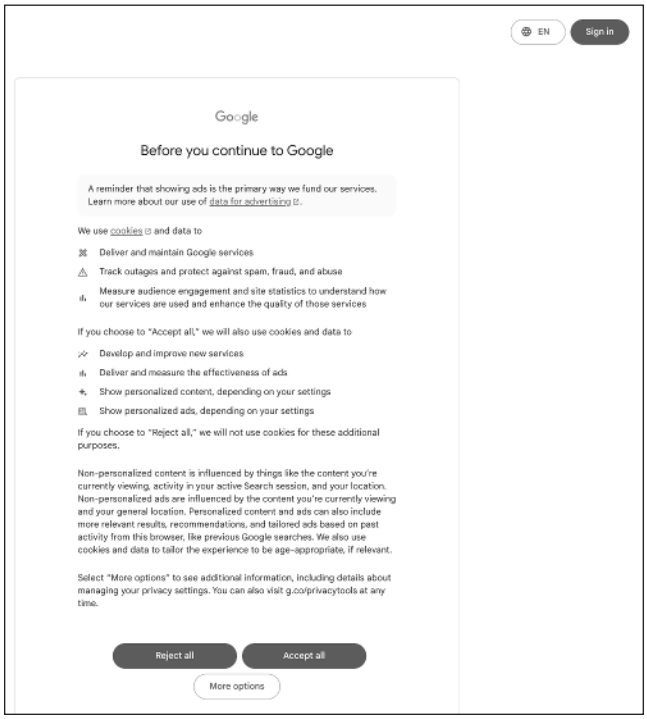


**FIGURE 1-1:** The notification that appears when you connect to Gemini (`gemini.google.com`) without first logging in to a Google account.

2. **Click either Reject All or Accept All.**

   Clicking Accept All means you're accepting Google's terms of use. Clicking Reject All means you don't accept cookies. You can also click More Options to define which permissions to grant Google regarding your preferences related to your personal data.

TIP

You can find all the information on Google's general privacy management at `policies.google.com` or by clicking the Privacy and Terms link, which usually appears at the bottom of the page of all services offered by Google. It is up to you whether to grant certain permissions or deny them in part or in full, according to your preferences.

Whatever you decide for permissions, the Gemini platform appears (see Figure 1-2).

3. **Click Sign In.**

Gemini asks you to enter your Google account login details, if you have them, or to create a new account.

TECHNICAL STUFF

If you do not yet have a Google account, creating one is easy. Once you click Create an Account, Google guides you through the process. In fact, it's so easy that I'm skipping along to logging in.

After you have logged in to your account (or created one if you did not have one), you see the same start page as the Gemini platform but with two differences: the Log In button, which has now become Try Gemini, and the circular logo that identifies your account, using your initials or a photo if you have chosen to provide one, in the upper right corner.

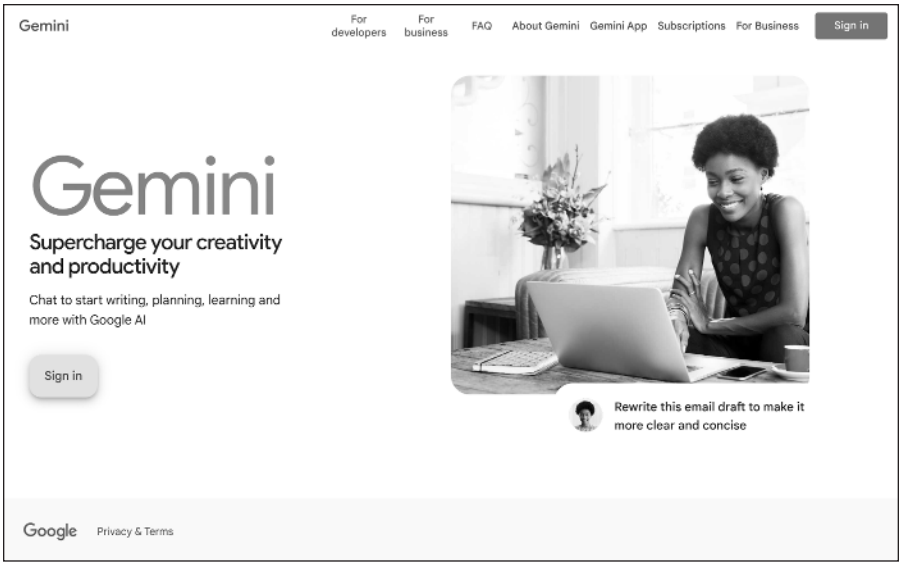


**FIGURE 1-2:** The Gemini introduction page, with the Sign In button shown at the top right and bottom left.

TIP

As you can see in Figure 1-3, Google provides a Help Center as well as guides for all its platforms, Gemini included. A Privacy Hub, related to privacy and personal data processing, is also available at all times (see Figure 1-4). You can (and actually should) consult this information before proceeding to use Gemini.

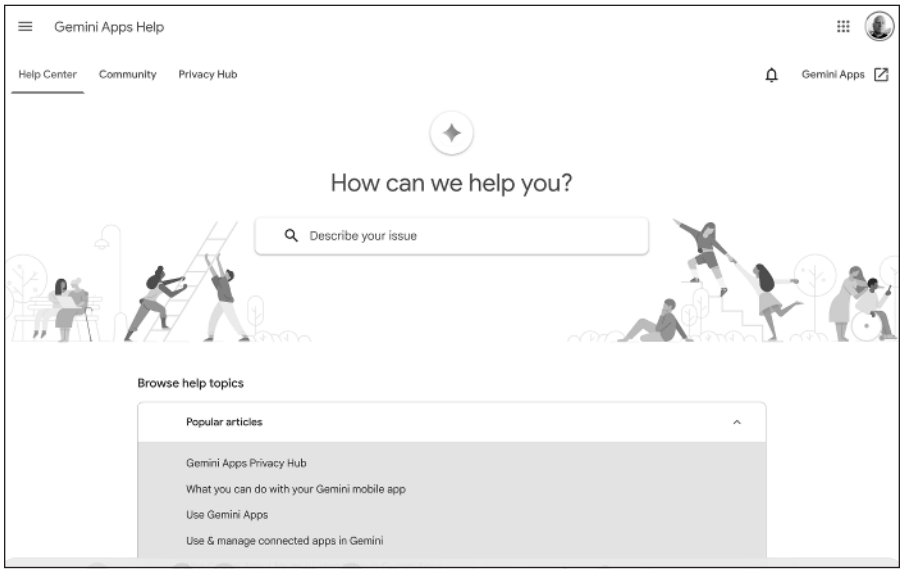


**FIGURE 1-3:** The Gemini guide, with all the basic information you need to get started using the platform.

**4. Click Try Gemini, or whatever invitation currently appears for first-time signup.**

Before giving you access to the AI platform, Google provides preliminary information, asking you to accept the terms of use once again. In this case, these terms concern your conversations with Gemini. By clicking Your Data Can Improve Google AI, you are provided information on how your data is usually handled. This information is always available on Google's Privacy Hub (see Figure 1-4). The subsection "Storing chat instructions" discusses how you can partially modify these settings, but for now, the only way to use Gemini is to accept the terms and conditions by clicking Try Gemini.

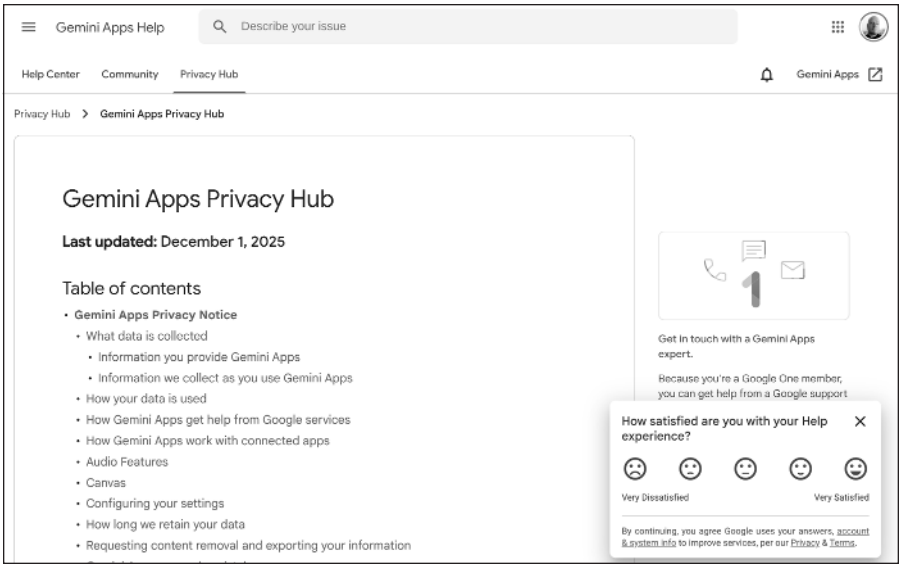


**FIGURE 1-4:** Google's Privacy Hub section related to all aspects of user data management in Gemini.

# Exploring the Interface and Main Options

After you sign in, Gemini presents you with the chatbot interface (as shown in Figure 1-5). You should note the difference between the chatbot and the platform:

- **Chatbot:** As the name suggests, an interface that allows you to dialogue with Gemini's AI
- **Platform:** The entire ecosystem surrounding the chatbot, and which you explore throughout this book

**REMEMBER**

As I mention in the book's introduction, content may have changed by the time you read this. To stay current on changes and to receive regular updates on the evolution of Gemini and Google's entire AI ecosystem, refer to the related section on the `aiskill.it` website (using your browser translation tools to get them in your language). To access exclusive prompt resources, visit this book's Cheat Sheet by going to `www.dummies.com`, searching for Google Gemini For Dummies, and then clicking the View Cheat Sheet link.

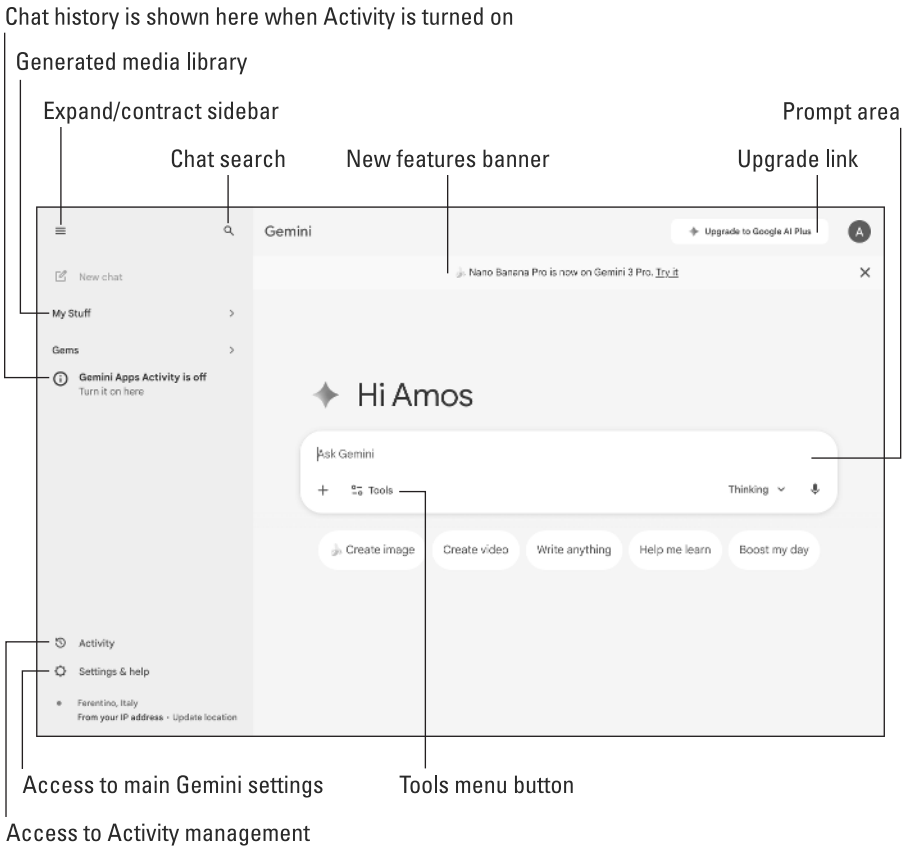


**FIGURE 1-5:** The Gemini interface, with its two main areas: the central Chat area and the platform area on the left with its various options.

# The chatbot interface

So that you can easily navigate as you practice with the various tools and options, the notable features of the chatbot interface are as follows:

- **Prompt area.** The field at the bottom of the larger central area, this is where you type instructions and requests for Gemini.
- **Side panel.** A collapsible area to the left of the prompt. You click the three horizontal lines at the top left to expand or collapse the panel. Notice that when the panel is in its collapsed state, only its main tools appear; but don't worry about them for now because you'll gradually explore the available features throughout the book.
- **Gemini.** Next to the side panel, you see the name "Gemini." This previously housed the available AI models; Google later moved this selection inside the prompt area.

TECHNICAL STUFF

Although you'll get used to these minor interface changes in time, it's important to notice them — and eventually learn to adapt them to your workflow — after grasping the basic concepts behind the graphic interface itself. AI models usually evolve over time, so are marked with a version number (for example, Gemini 2.5, Gemini 3, and so on) as well as a name that identifies their "power" (Flash, Pro, Fast, Thinking, and so on). You can quickly obtain more details about the AI models by querying Gemini.

- **New features.** At the top center, Gemini displays notifications about new features whenever they become available. At the end of 2025, for example, these messages related to the availability of the tool initially known as Nano Banana or the introduction of the Gemini 3 model. So don't be surprised if you see similar messages when you log in. You can always click the X to the right of each message or the Accept button to get rid of it.
- **Upgrade button.** In the upper right corner, next to the circular icon identifying your Google account, this button enables you to activate a subscription if you use the free version of Gemini. If you've already subscribed to one of Google AI's plans ("Choosing Between the Free and Pro Versions" later in this chapter) you'll see an Invite a Friend button to share a free initial subscription to premium packages.

## The prompt area

Before you send your initial request to Gemini, you should focus on the prompt field (see Figure 1-6). Quickly explore its tools before sending your initial request to Gemini.

- **Tools.** At the bottom of the main prompt area, you'll see a series of items representing Gemini's currently available tools (which I discuss one by one in this chapter.

  These tools previously appeared below the prompt area as buttons. Google increased their number over time, so you now access them from a single button inside the prompt area itself to display a vertical menu — like the one that was already standard for attachments (another example of typical changes to the interface).

TIP

- **Plus sign.** On the left, a plus sign allows you to add something to the instructions typed inside the prompt area, such as attaching one or more files to your request — which, as you'll see in practice in Chapter 4, greatly expands the possibilities for AI queries.

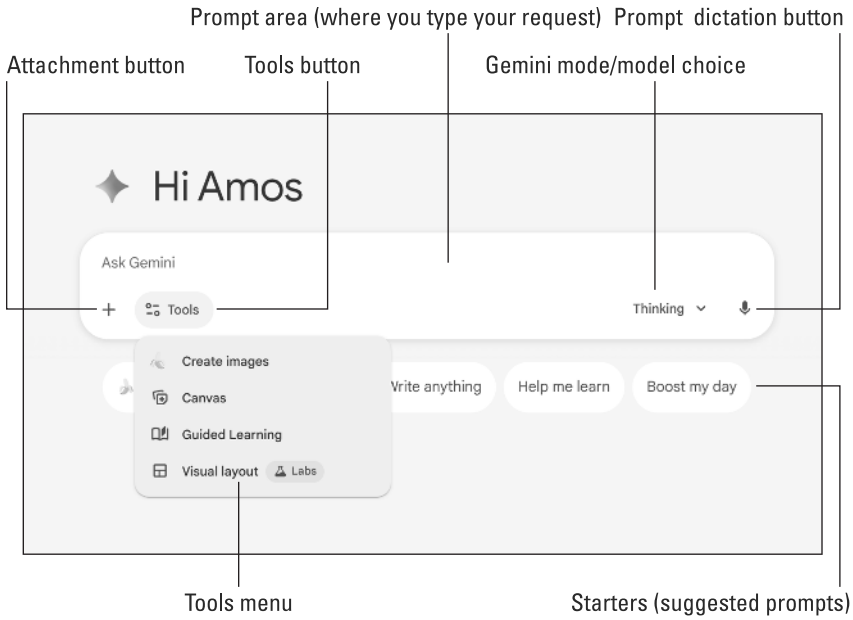


**FIGURE 1-6:** The prompt field, the area where you type instructions for Gemini, with its various tools.

TECHNICAL STUFF

In the field of AI, the ability to recognize, as well as generate, content that goes beyond simple text is called *multimodality*. Although absent in the initial launch phase of Gemini, this feature was subsequently introduced and enhanced to offer users ever greater flexibility and power. Today, you can add both multimedia files (images, audio, video) and documents in various formats (text, spreadsheets, PDF, and so on) as attachments to the prompt.

## Let Gemini describe itself

Initially, in my books about AI, I used the first few chapters to introduce theoretical concepts as well as technical and historical information. This kind of information is becoming increasingly redundant because you can obtain it from AI itself. To see how this works, follow these steps:

1. **Type the following instructions into the prompt field (see Figure 1-7).**

   ```
   List the entire history of Google's AI from Bard to today's models inside
   a table, mentioning all the tools and models available over time.
   ```

   As you type, notice that an airplane symbol appears in the lower right corner of the prompt area, replacing the microphone symbol. ***Note:*** You use the microphone symbol to dictate instructions by voice, so it's useless when you type.

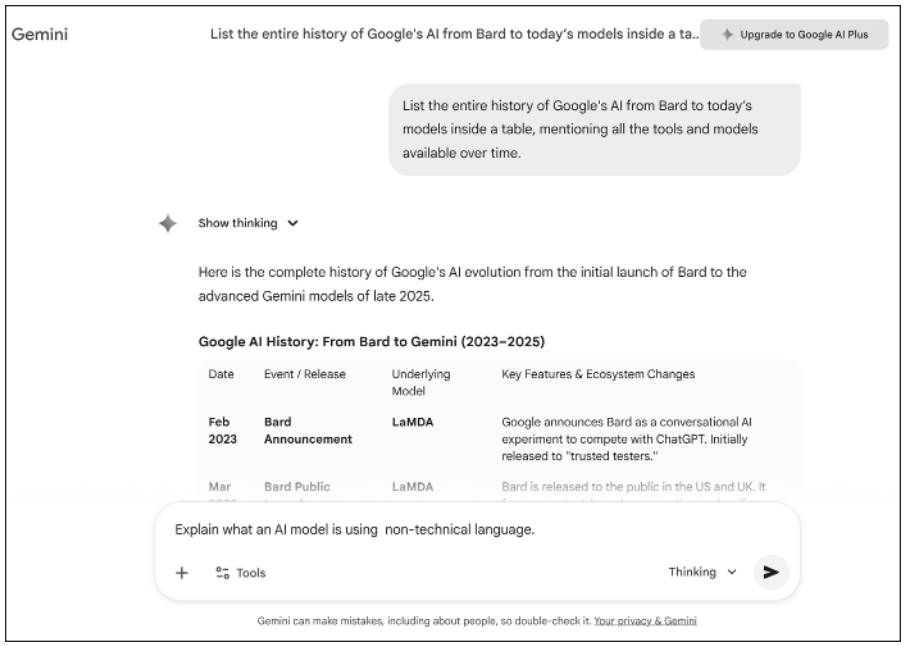


**FIGURE 1-7:** The initial part of the chronological table obtained with your first prompt and, below, a new prompt that extends the chat.

2. **Click the airplane symbol or simply press Enter.**

TIP

Pressing Enter doesn't break the text and create a new line, so when you need to go to the next line in a prompt (for example, to create a list) hold down the Shift key while pressing Enter to prevent instructions from sending before you finish typing them.

After you send your request, Gemini should include a table in its response as you requested. In this case, the Export to Sheets option appears near the top of the table (see Figure 1-8).

If you already use Google Workspace tools — for example, the cloud applications the company offers — you know that the Sheets app allows you to create spreadsheets, just like Microsoft Excel. If, on the other hand, you are not yet familiar with these tools, clicking Export to Sheets ensures that the table transfers to a Sheets document. You can then open the spreadsheet containing the information in the table using the appropriate button that appears at the end of the operation (examples of Google Workspace are discussed in Chapter 9).

REMEMBER

This is an example of a basic prompt. Chapter 4 discusses in detail the techniques and strategies for composing structured and effective prompts, creating a tool that you can use both to learn this AI query technique and to generate prompts starting from a simple request.

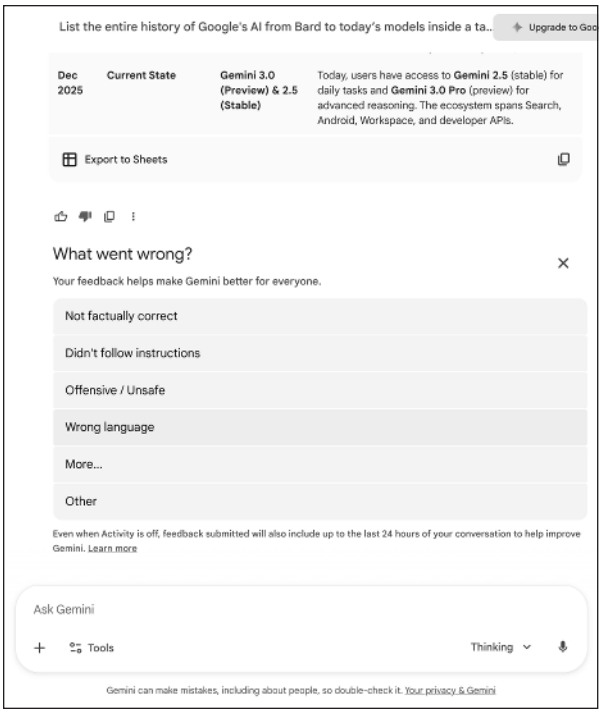


**FIGURE 1-8:** Negative feedback (thumbs down) on an answer and the various detail options available. Also note the menu that appears when you click the three dots.

# Managing Gemini's responses

When you move the mouse pointer over a response's content, a series of buttons appears, followed by three vertical dots; each icon displays an identification label when you hover over it. In this section, you'll see how to use them to improve the chatbot user experience.

## Good and Bad response

The first two buttons should be familiar to you — the like/dislike icons are found everywhere. You can click them to register, respectively, a Good or a Bad response, so Google can improve its AI based on your feedback.

When you provide positive feedback, the system displays a thank-you message to acknowledge it. Selecting "thumbs down" means you'll be asked for more details about your negative rating (refer to Figure 1-8). Clicking More expands the list of options shown, whereas clicking Other opens a field where you can type details of the problem you encountered. If you change your mind, rather than sending your feedback, you can click the X in the upper right corner and close the list, then click the thumbs-up button instead.

## The Redo option

The two buttons just described are followed by a circular arrow, the Redo button (see Figure 1-9). You can use this button to request a new response.

TECHNICAL STUFF

Due to its training, generative AI gives responses based on statistical and mathematical criteria, each time introducing an element of randomness. This results in responses that are always different, even when you make the same request. To explore this aspect further, simply enter the first two paragraphs of this note in a prompt, enclosed in quotation marks and preceded by "Explain this concept to me in simple terms." For example, you can do this after asking Gemini to explain what an AI model is, as shown back in Figure 1-7, or for any technical concept you may find unclear.

After clicking Redo, two navigation arrows appear on either side with numbers in the center (for example, < 2/2 >), which you use to browse all generated responses (also shown in Figure 1-9). These arrows, like the like/dislike buttons, appear only for the last response as you continue chatting.

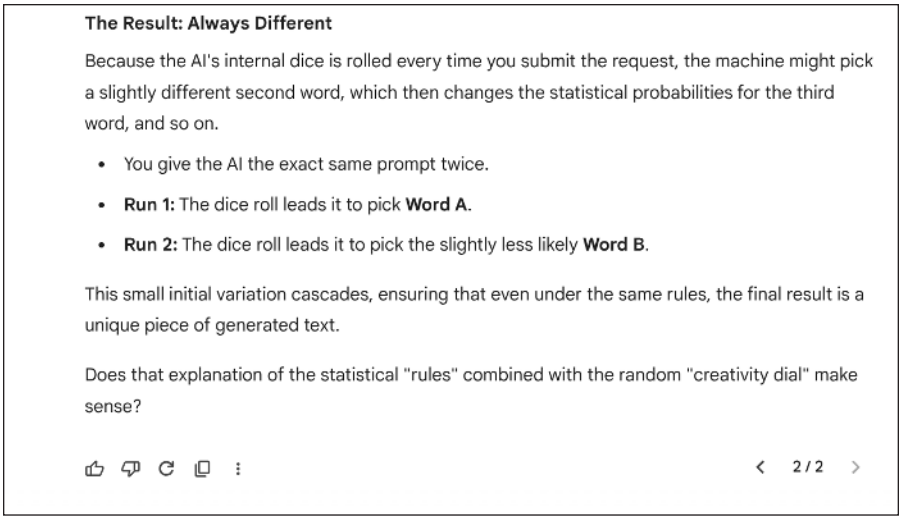


**FIGURE 1-9:** The two navigation buttons (bottom right of the image) that appear when you regenerate Gemini's response using the Redo button.

## The Sharing option

Recently this option moved from the Chat area to the right corner of the screen, and it may move again. That said, chats will always have a standard sharing button. When you click this button, you create a link to share, either privately or publicly (see Figure 1-10). At the time of this writing, you can also copy and paste a link to the chat or post it on various social networks (see Figure 1-11), allowing others to view it simply as any web page or repost it in their own account. You also

have an option to export the chat content to Google Docs; for those unfamiliar with it, this is another Workspace tool, comparable to Microsoft Word. You can also export to a draft message in the Gmail email service. Again, these options may not be fully or partly available in some countries.

WARNING

Remember to always share links to your chats responsibly.

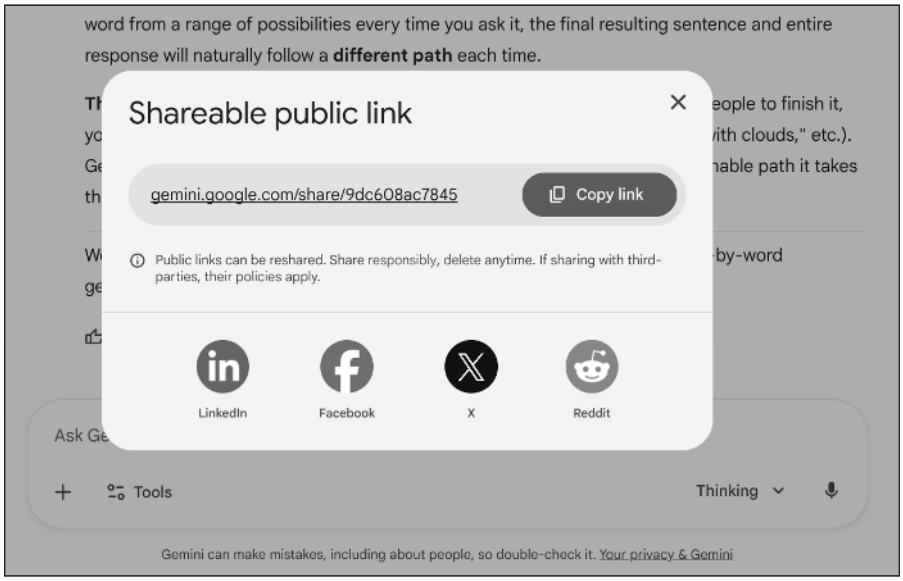


**FIGURE 1-10:** Conversation sharing options available in Gemini through the typical icon, usually shown in the top right corner and for each chat listed in the left panel.

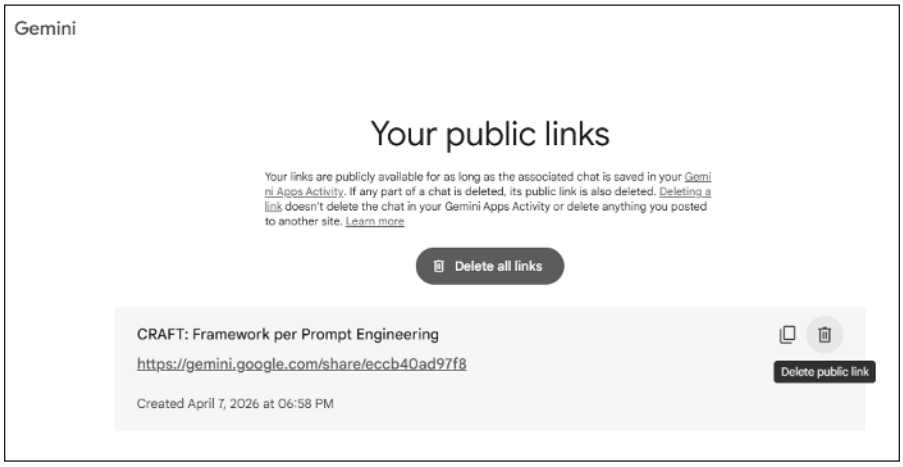


**FIGURE 1-11:** Your Public Links screen, listing all the links you created, with options to copy or delete them.

## Copy Response and more options

Next to the Share option is the button that allows you to copy the content of a single reply and paste it into any text document.

Clicking the three dots next to it gives you access to several more options.

The first option starts a Google Search which allows Gemini to verify the accuracy of the information provided in the response (see Figure 1-12). If the search is successful, verified information is highlighted in green and followed by a small downward arrow, which you can use to find out the source queried and provide feedback on it if necessary.

If, on the other hand, Gemini finds inaccuracies (for example, due to a "hallucination"; see the next warning), text is highlighted in red and followed by notes citing sources that refute that specific piece of information.

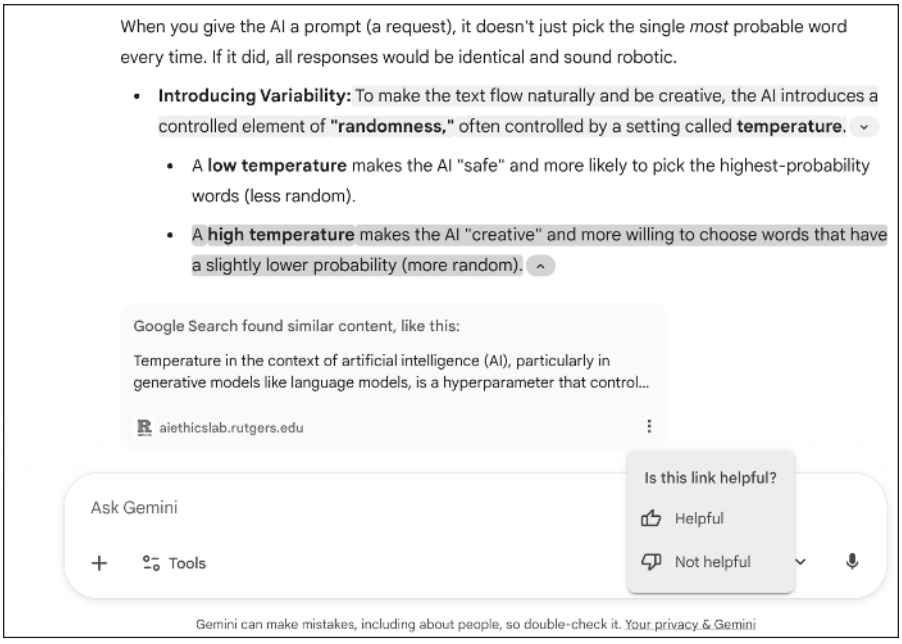


**FIGURE 1-12:** Parts of the response highlighted after verification and details of the source used for verification.

WARNING

As many figures in this chapter show, a message appears below the prompt field warning that Gemini may make mistakes in its responses. Again, this is because of the intrinsic mechanism by which generative AI works due to its training process and algorithm. These errors are referred to as *hallucinations* when they appear in responses (this term can also be explored in more depth by asking, as usual, Gemini itself to explain it in simple terms). In addition, responses may contain outdated information, despite Gemini's ability to conduct online research. During tests conducted for this chapter, for example, Gemini included its new model 3 only once in the final part of the chronology generated by the prompt (refer to Figure 1-7 and Figure 1-8), albeit being available in the chatbot itself.

Other available options may allow you to listen to the response with a synthesized voice, report a legal issue (for example, the production of dangerous or defamatory information, or information that may have infringed a copyright), and the previously mentioned options for exporting the text to Docs or a Gmail draft. If legal issues are found, a special Google page appears where you can provide the relevant information and send the report.

## Chat memory and "system memory"

If you decide to continue the conversation as suggested, the AI keeps track of what has already been discussed. Each chat is a kind of container (actually a sort of web page) in which you can refer to what was mentioned previously, without having to provide details again, just as happens during a normal conversation between humans. The *memory* of a chat — that is, the space available to continue with requests and responses — is not unlimited, but Gemini's is quite large. The chatbot will eventually warn you when it exceeds the limits, beyond which the AI begins to "forget" the initial part of the conversation.

You can, of course, also send your requests in separate chats, which is certainly a more astute choice when dealing with different and unrelated topics.

WARNING

When you start using the Gemini chatbot, check that the Keep Activity status in the left panel is turned on in case you need to save your chat. You may have previously disabled chat history, for example, to prevent Google from storing location history or other activities related to your account. If activity tracking is disabled, clicking New Chat triggers a warning: "Gemini app activity is disabled." Clicking the Enable link below this warning opens a window where you can toggle history on or off (see Figure 1-13). Clicking Delete allows you to delete all your history within a certain timeframe. You can also browse and manage all past activities in your account using the list shown on the rest of the page. Of course, you can return to

this screen to manage the retention times of your activity history and selectively or completely delete them as needed.

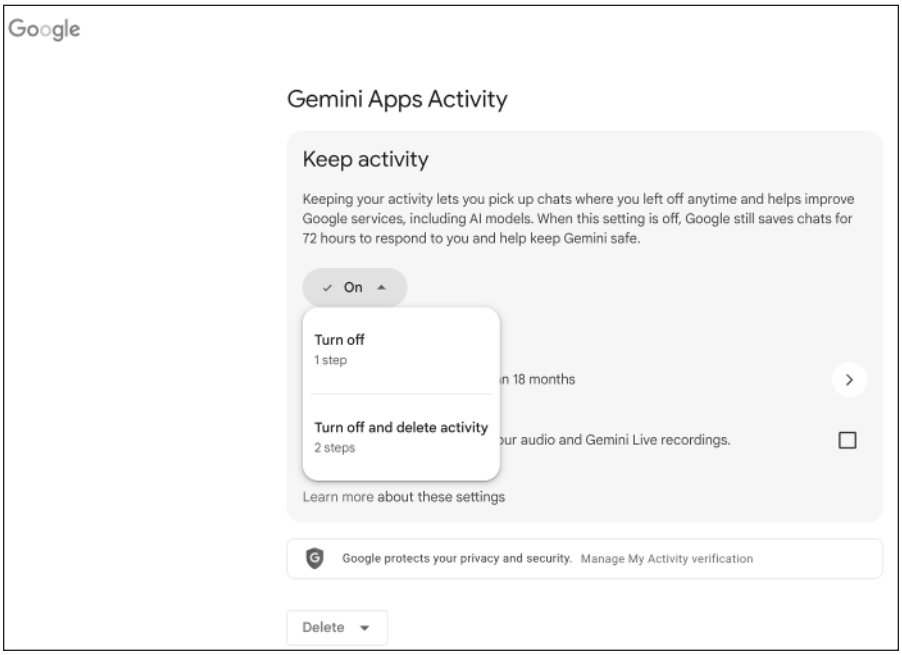


**FIGURE 1-13:** The Google page that appears when you decide to activate Gemini's history (if it's not already).

## Chat history and temporary chat options

Once you activate chat history, each new chat adds to that history. The Gemini chatbot displays it all in the left side panel (see Figure 1-14). Gemini automatically assigns a name to the stored conversation based on its content, but you can rename it for easier identification. The renaming option — along with those for sharing the entire chat, pinning it to the top, and deleting it — is available in the menu that appears when you click the three dots to the left of the chat name (also shown in Figure 1-14).

**TIP**

Next to New Chat, notice a button showing a dotted speech bubble. You can click this and start a temporary chat — that is, one whose content is deleted once you start the next chat. This option is particularly useful when you want to immediately delete all traces of a conversation containing sensitive data.

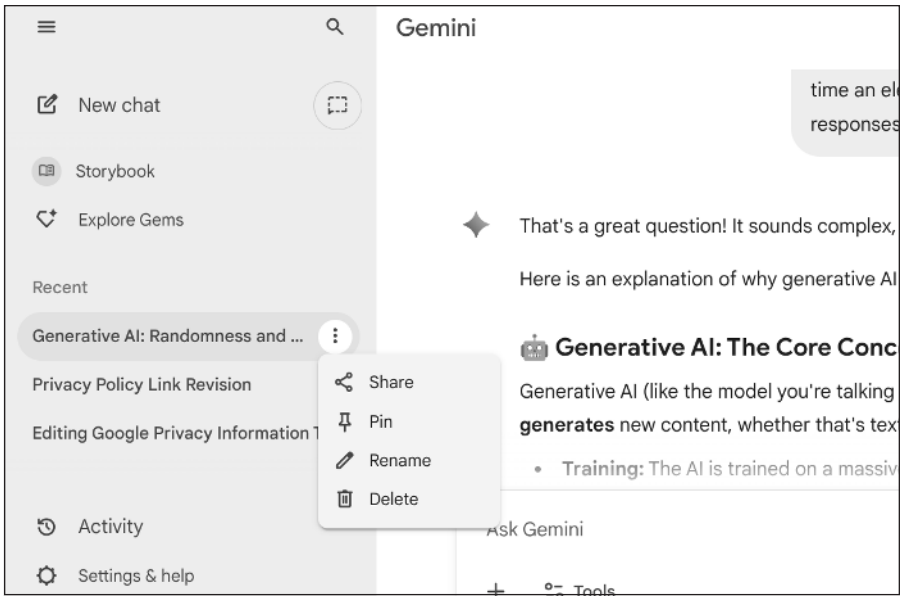


**FIGURE 1-14:** The Gemini chatbot side panel with a partial chat history and, next to the most recent one, the menu showing all available options.

## Storing chat instructions

More important than the ability to delete information from Gemini's memory, you can also store instructions for the chatbot to remember and apply to all your conversations. To access, activate, and configure this important feature

1. **Click Settings & Help at the bottom of the left panel and select Instructions for Gemini (see Figure 1-15).**

   A new screen appears with a button for adding these instructions, and an "active" switch in the upper right corner allows you to turn instructions on and off (see Figure 1-16).

   The page also contains an extra button that shows some example instructions and a brief introductory text. As previously stated, these options may be slightly different when you read these pages, so just keep in mind the general directions just described as well as the option names or positions.

   Clicking the Show Examples button simply displays a floating panel with a list of sample instructions.

2. **Click Add.**

   A window opens in which you can type your personalized instructions for the chatbot; these are then stored, allowing Gemini to follow them during your new conversations.

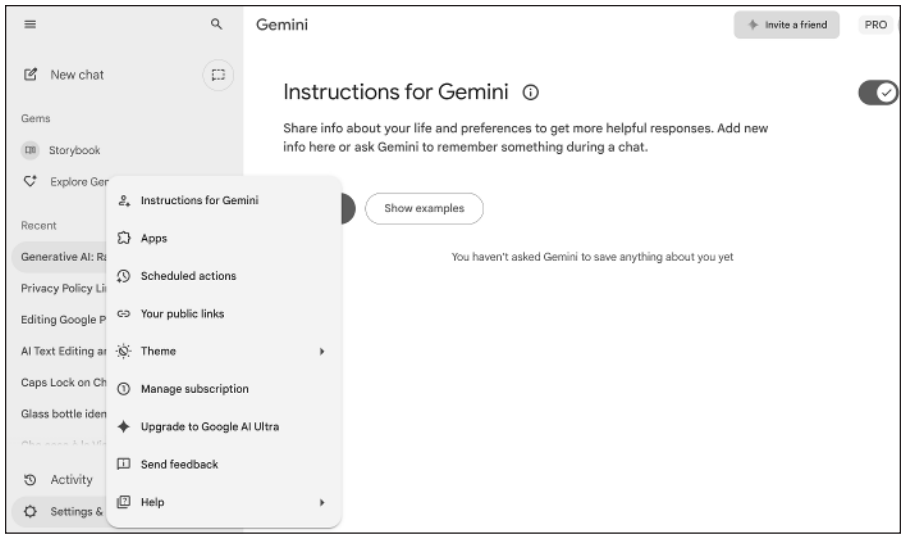


**FIGURE 1-15:** The settings menu, which you can access from the menu item of the same name at the bottom of the side panel.

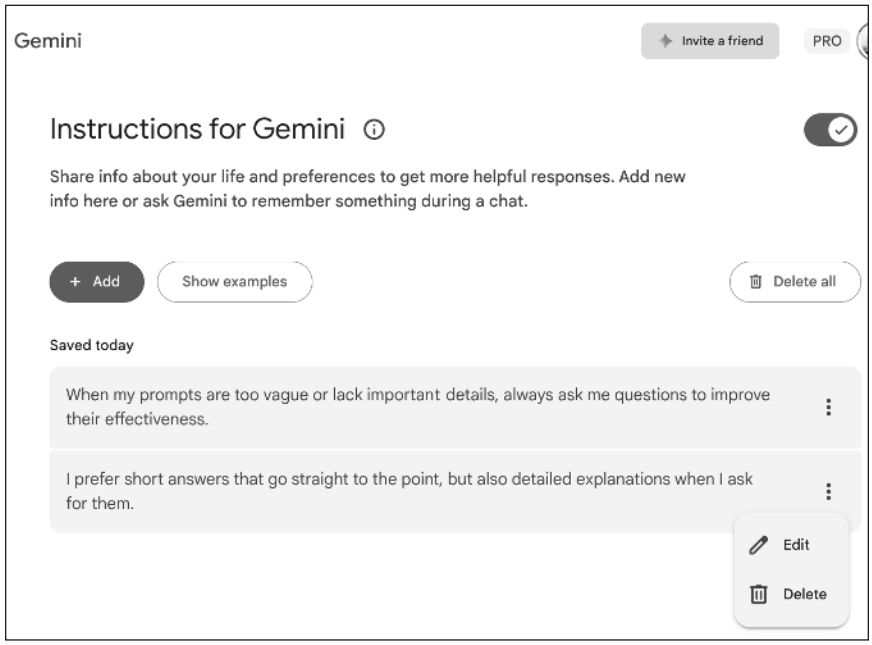


**FIGURE 1-16:** A few personalized instructions and the menu to edit or delete them individually.

To better manage instructions, referred to as *system instructions* because they apply to all activities, you may find it useful to group them into separate items. This allows you to edit or delete each individual item at any time. To change or delete them, click the three vertical dots to the right of the desired item (refer to Figure 1-16) and either click Edit or Delete.

REMEMBER

At the time of writing, Google just included a "chat memory" in Gemini. You may notice it when the chatbot refers to information you have mentioned in past conversations. Other platforms (the first and most efficient in this area being ChatGPT), have already extended this feature and made it an important aspect of chatbot personalization, tailoring it to your AI needs and habits. Hopefully Google will do the same for its Gemini in the future. Until then, you can update the Instructions for Gemini section just by telling the chatbot to "remember" something in the chat itself (you can find these "memories" in the related page as shown in Figure 1-16).

# Choosing Between the Free and Pro Versions

Like all other companies offering AI freemium platforms, Google attempts to offer most of its available features and tools in its AI ecosystem to you for free, although in a limited form. However, the cost of maintaining the related infrastructures is high, creating the need to limit your access to certain features (usually hourly or daily) if you don't have a paid subscription.

Google offers a variety of subscription plans (as shown in Figure 1-17). Features differ from the free plan and include cloud storage space and credits used for generating video content with Veo, the AI video generation model that I discuss in the next section.

The Ultra subscription, the most expensive of all and shown on the right in Figure 1-17, provides the most power and space, but it is clearly intended for business or highly professional users.

REMEMBER

The costs of subscription packages usually change over time, as do their features. Initial offers tend to vary as well — between a free trial month and two months at a reduced price. In some cases, subscribers are offered an invitation coupon to share with a limited number of people, allowing them to register with Gemini and enjoy one or more free months. If you decide to use these offers just to try out the service, remember to cancel your subscription before the first charge — using a reminder, if necessary — because costs are automatically charged to whatever payment system you set up to take advantage of the offer.

Finally, as shown in Figure 1-18, the different subscription tiers (excluding the cheaper Plus plan) are no longer linked to the number of credits you can use to generate video content — now 1,000 for all plans except the very expensive Ultra — whereas cloud storage space (on top of the list) varies considerably among all tiers.

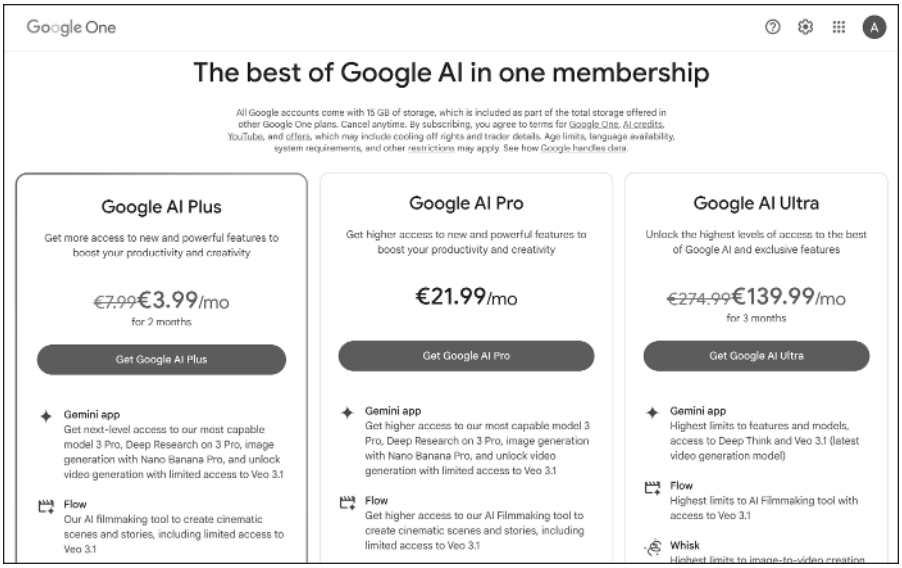


**FIGURE 1-17:** Google AI subscription packages available at this book's writing. They may change, so always consult the official page for updated prices and features.


**FIGURE 1-18:** Details of the subscription plans from Plus to Ultra showing the number of credits and storage space.

One feature not mentioned in the list of different packages is so-called Gems (see Figure 1-19), a kind of personalized chatbot, which I discuss in Chapter 4. Google initially made some Gems available to handle specific activities such as brainstorming, text review, and more. The most interesting one, introduced in the summer of 2025, is Storybook, which allows you to generate entire illustrated books from a prompt (discussed in Chapter 5).

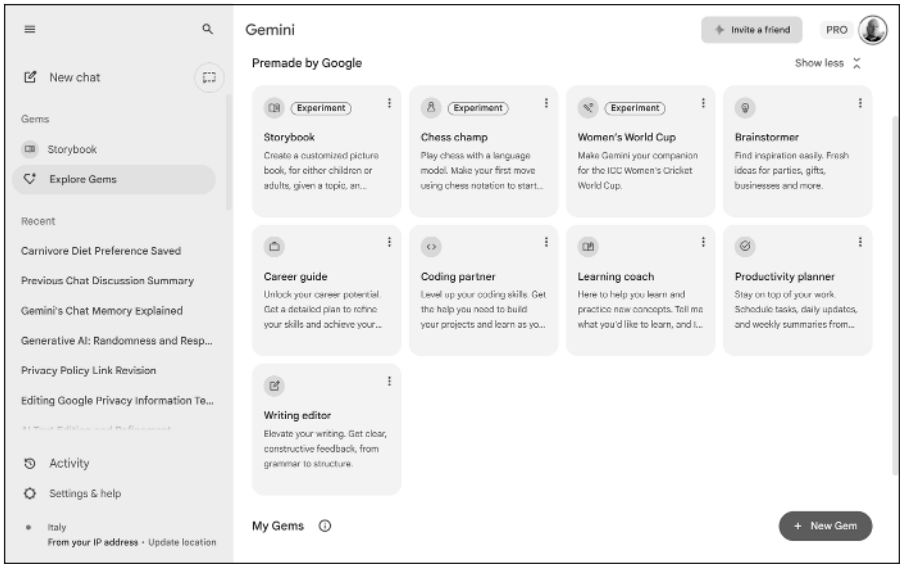


**FIGURE 1-19:** Several Gems available from Google, accessible from the Explore Gems option in the left panel; the button for creating custom Gems is clearly visible on the right.

TIP

Because they're free, users often register more than one Google account, thus taking advantage of AI use, because you can use Gemini on multiple accounts and increase access to limited features in the free tier. Google's entry subscription plan, called Plus, often includes a free initial month or two at an even cheaper price, thus allowing users to try some of the premium features before deciding if this or even more expensive plans are worth the money. In any case, the goal of *freemium* philosophy (a cross between free and *premium*, paid) is precisely to convince users to embrace paid packages after having "tasted" the features on offer.

TECHNICAL STUFF

At the time of writing, Google's most powerful free AI model (currently Gemini 3) includes daily prompt limits, which, once met, generate a warning informing you that you can no longer submit prompts for a set period. Given this, consider doing the following:

- Use the most powerful model for creating complex content, such as the meta-prompts discussed in Chapter 4.
- Use basic models for simpler tasks, such as applying meta-prompts in new chats. At present, these models are more than capable of handling such tasks.
- Consider a paid subscription for professional work. This reduces prompt limitations and eliminates the need to switch frequently between models, even though daily usage will still not be unlimited.

# Discovering Multimodality

A hallmark of the digital age, *multimodality*, is the practice of combining multiple forms of communication that can include video, audio, images, and text. Its purpose is to help people interact and create meaning — something that Google Gemini offers in spades.

As I mention in the previous section, activating a subscription unlocks certain features, whereas others already available in the free version are extended in their use. For example, the image editing feature (see Figure 1-20) or the use of the more powerful model Gemini 3 (both introduced in 2025 and rapidly enhanced before the end of that year) are just two features available without a subscription, but with limitations on the number of requests that you can send daily.

The image editing feature, nicknamed Nano Banana, actually evolved from the Gemini 2.5 Flash model, with the term "image" added to define its role. This model improves Gemini's image generation capabilities by adding the ability to modify images. For example, you can mix elements from different images into a new image or edit details of the original image while preserving its original features, thus improving consistency.

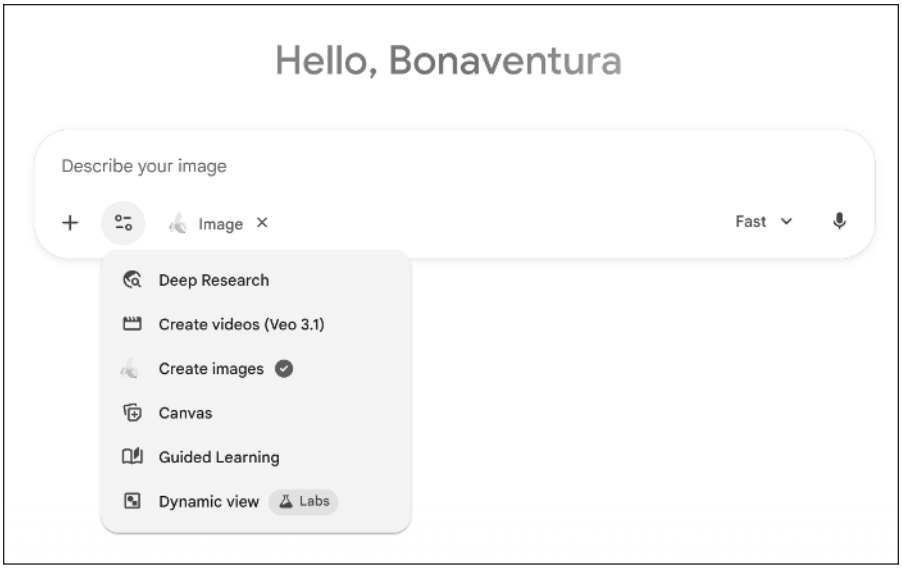


**FIGURE 1-20:** The prompt field with the Nano Banana image tool active. Note the presence of an experimental tool at the bottom of the menu.

Gemini's multimodal AI can manage multimedia content in two ways:

- **Inbound:** By attaching it to the prompt or capturing it live
- **Outbound:** Generating it from scratch using text descriptions and/or based on attached files

  For example, by attaching an image and using the necessary instructions, you can analyze the image's characteristics and use them to generate similar images. You can also extract information from a complex PDF file containing both text and images, or even analyze video and audio files.

You'll explore features like these in Chapter 4. As for Gemini's generation features, they are as follows:

- **Image generation.** In November 2025, Nano Banana further evolved to its Pro version, with an even greater consistency between generated images containing the same subject and other general graphics improvements. By the time this book reaches you, the model will have surely further improved, along with all other Google models and tools.
- **Video generation.** Another evolutionary leap in Gemini's multimodal capabilities, video generation is obtainable both from images and through simple instructions (see Figure 1-21), with the third version (Veo 3) having been released in the summer of 2025 and updated further to 3.1 in the fall. This new generative model for video, available only to subscribers at the time of writing, has blown away the competition by the introduction of an audio track.

  
  REMEMBER

  Generated videos are managed using credits, which, as I note earlier, initially vary in quantity depending on the type of subscription (refer to Figure 1-18). For example, the basic subscription plan allows the creation of just two videos per day. Video duration and resolution are also linked to the type of subscription. Unfortunately, these limits are not clearly indicated in the options listed in the different types of subscription, at least at the time of writing.
- **Audio generation.** As you'll see in the dedicated insights, audio — or rather, speech synthesis — plays another important role in Gemini's multimodal features and in NotebookLM (another platform that uses the same AI and discussed in Chapter 2). For example, you can generate a sort of podcast, called Audio Overview, in both platforms from attached or generated content, using a command that appears after file processing (see Figure 1-22). Although unavailable in the Gemini free or Plus plan, you can access it in NotebookLM even in the free version. Conversely, you can analyze and transcribe (or synthesize) an audio file by attaching it and then requesting its processing with a prompt. For speech synthesis, you can use your computer's microphone or mobile device to "talk" with Gemini's AI in real time — an

aspect explored in more detail when I discuss the mobile app at the end of this chapter. You can also dictate your instructions instead of typing them, or "listen" to the answers obtained in a text chat. In all these cases, Gemini transcribes the audio of the voice conversation for you to consult at your convenience.

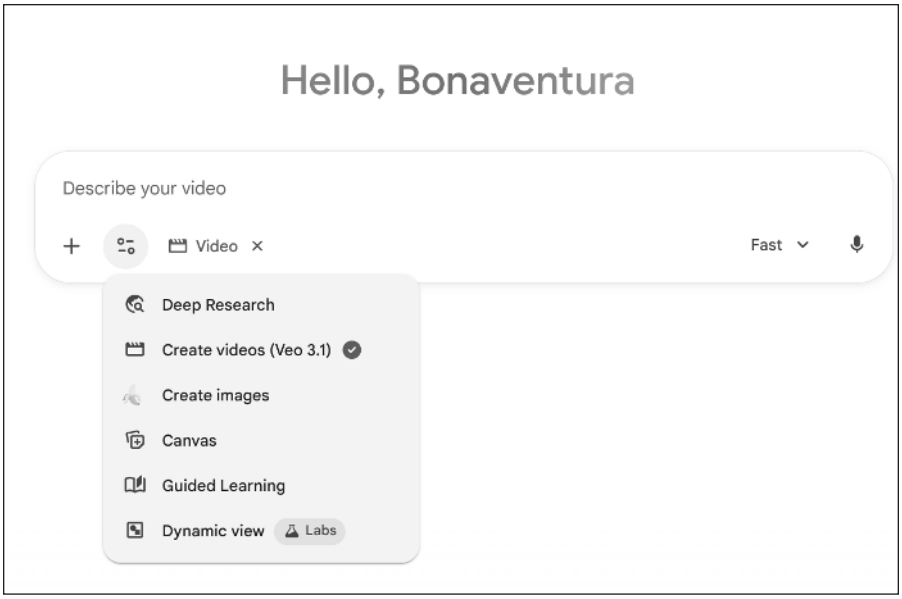


**FIGURE 1-21:** The video generation tool available in the paid version of Gemini, based on Veo.

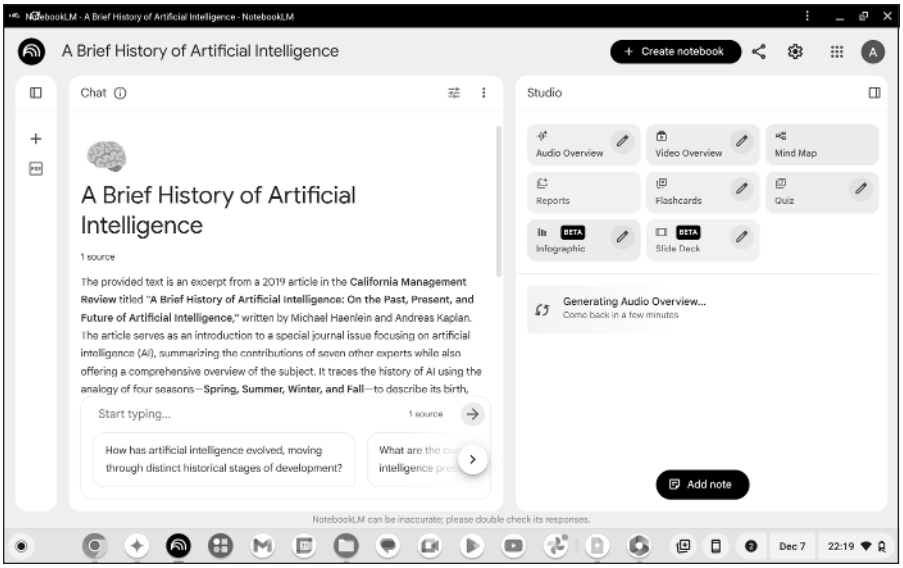


**FIGURE 1-22:** The generation of an Audio Overview from a PDF document in the free version of NotebookLM. Note the Audio Overview tool, listed first in the Studio panel on the right.

WARNING

The creation and use of multimedia content is a sensitive matter; you must not infringe upon copyrights, ethical standards, or legal rights — including privacy. Gemini often displays a warning to this effect before allowing you to use corresponding tools. In some cases, it prevents you generating multimedia content if it identifies elements in the prompt (or in the source file) that may violate conditions for the use of the tools on the platform, such as descriptions of sex or violence.

REMEMBER

Gemini supports attachments and conversations in many languages. However, while being inherently multilingual, most generative AI models are primarily trained on English datasets. Consequently, performance is often superior in English, thus allowing mother tongue or native English speakers to make the most of this technology. Also, it's important to know that even when using other languages, English remains the central internal bridge for translation and processing.

# Exploring Gemini and Google Applications

Google manages and makes available to you an entire ecosystem of applications (both on desktop and mobile devices). It was just a matter of time before its AI technology gradually influenced and eventually became part of the apps.

Clicking Apps from the settings that appear from the bottom of the left sidebar of the chatbot, you can access the secondary address `gemini.google.com/app`, which opens a welcome banner (see Figure 1-23) inviting you to explore this ecosystem interconnected with Google's AI.

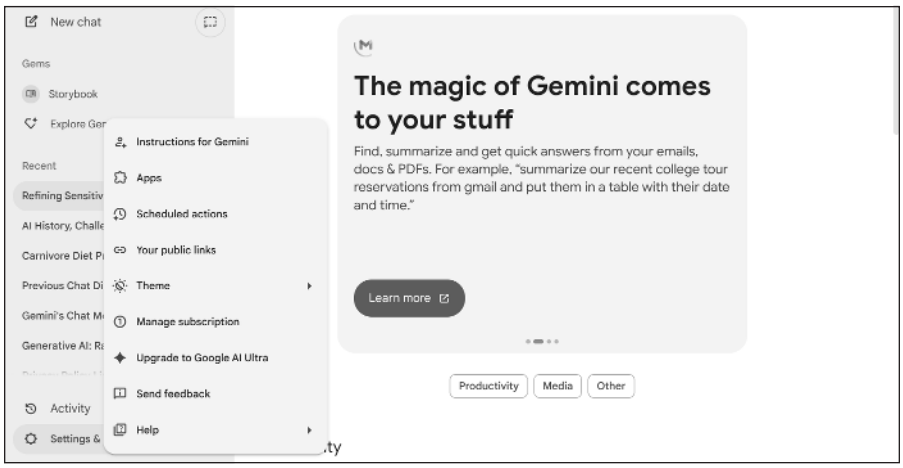


**FIGURE 1-23:** First access to Google applications connected to Gemini.

## Available apps

These applications (referred to simply as *apps* in this context) should already be quite familiar to you, especially Gmail, Calendar, and Workspace. You'll also find YouTube (including the Music version) and other minor but popular apps dedicated to travel and tourism such as Google Maps, Hotels, and Flights.

You can also see these applications are divided into categories on the settings page: Productivity, Media, and Other (refer to Figure 1-23).

REMEMBER

If you want to access settings for Google apps (including their connection to Gemini), you'll find you can only fully manage them using Chrome, the company's official browser. In other browsers, including Edge, they may not be fully available. Furthermore, during the first draft of this chapter, Google made the Travel item available, but later removed it; but you may still find it when you access these Google apps settings. As I continue to point out, changes in options and tools are inevitable during this AI's ongoing evolution.

## Enabling apps

Although some apps are already connected to Gemini by default, you can only enable others using a toggle (shown on the upper right of Figure 1-24). Upon first activation, as usual, a temporary confirmation window appears, providing basic information and a link to the page explaining how personal data is managed (which, as you discover Chapter 12, is often used to train AI).

The possibilities for using Gemini in Google apps are virtually unlimited, but it is easy to get an overview (as shown in Figure 1-25) by prompting Gemini with a request such as:

```
Provide an overview of how Gemini integrates with the Google ecosystem,
including Workspace, Gmail, YouTube, and Maps. Please summarize the potential
use cases and functionalities in a table.
```

You can also invoke apps directly in the prompt using the @ symbol (see Figure 1-25), combining them with the instructions in your request, which in this case are executed by using the selected applications. Depending on the apps available, you can, for example, call up YouTube as a source for a search, Maps to find accommodation or restaurants, Calendar to check availability of free times and dates and enter appointments, and so on. Results may not always be optimal, and in some cases your requests may even return an error of unavailability, but the quality of responses is often related to the search context. Furthermore, in Gemini's free plan, the use of this tool is also limited, as it is with other advanced features.

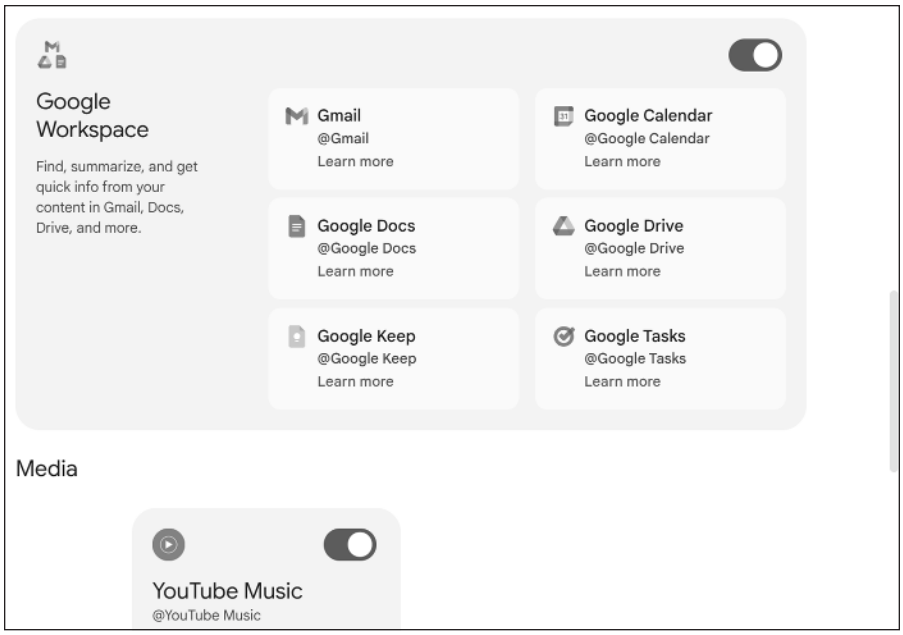


**FIGURE 1-24:** An example of apps in the Productivity and Media categories that you can connect to Gemini, the former mainly from Google Workspace.

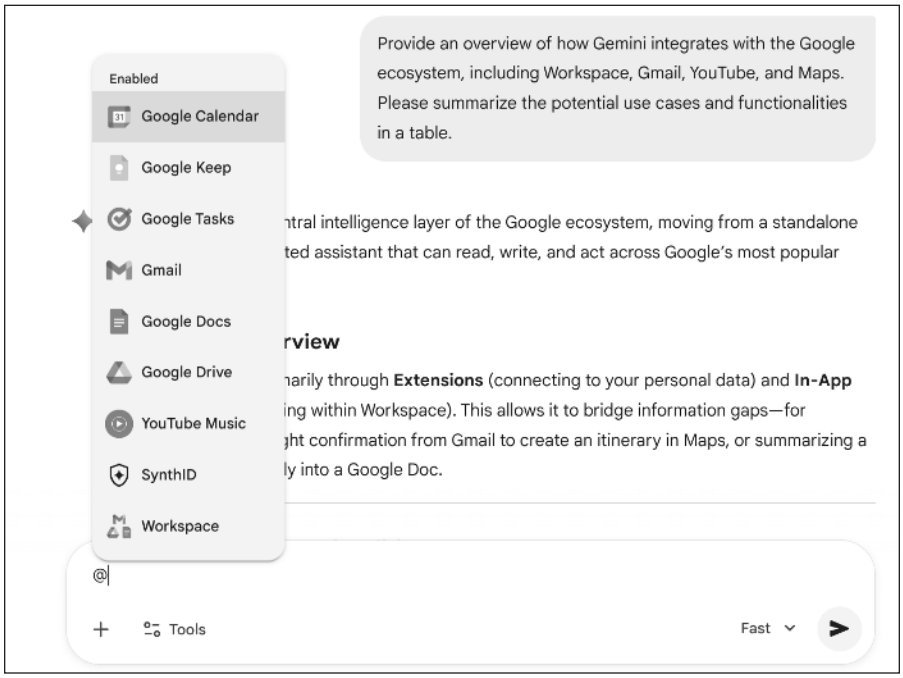


**FIGURE 1-25:** The prompt that produced an overview of Gemini's use in Google apps. Note how the @ character "calls" the apps when used in the prompt field.

REMEMBER

In Part 3, you'll use Gemini in conjunction with various Google apps, executing tasks that involve NotebookLM and other AI tools. Practicing with different tools provides a more comprehensive overview of Google AI ecosystem's practical use. An entire chapter would probably not be enough to describe the countless possibilities for using Gemini in the context of Google apps, and a complete guide would probably require an entire book. However, as you will see in Part 3, this tool is extremely intuitive to use and, as you will discover in Chapter 3, you can get help from Gemini's AI by sharing your screen with the chatbot and talking with it in real time about any computer activity. This also applies to the Gemini app on mobile devices, discussed in the next section.

# The Gemini Mobile App

In addition to the web chatbot, you can also use Gemini as a mobile app on smartphones and tablets. You download the app from the Apple store (see Figure 1-26) or Android store, based on your smartphone or tablet operating system.

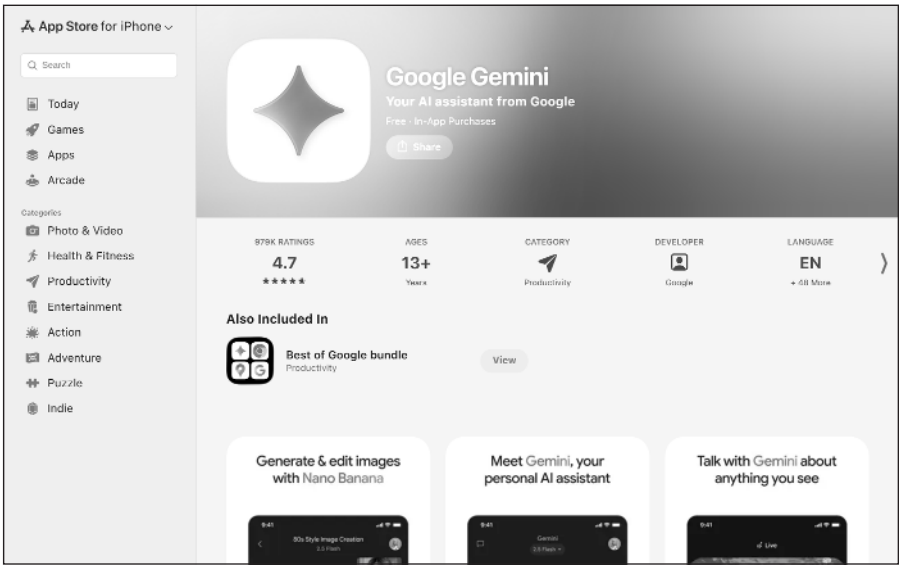


**FIGURE 1-26:** The Gemini app in the Apple store; note the large number of ratings and the positive score.

REMEMBER

Of course, features and options in the mobile version of Gemini (shown in Figure 1-27) may differ from the desktop version — that is, the one accessible from the web. Still, the app offers some features specifically designed for use on mobile devices, thus making it useful for certain tasks, especially those using the device's camera. You'll learn more about them in Chapter 4.

As shown in the three screenshots in Figure 1-27, you display the left side panel containing the chat history and other options (center image) by tapping the three horizontal lines in the upper left corner (left and right images). Meanwhile, you access available options for querying the chatbot by tapping the Tools button next to the + sign at the bottom left of the prompt field (left image).

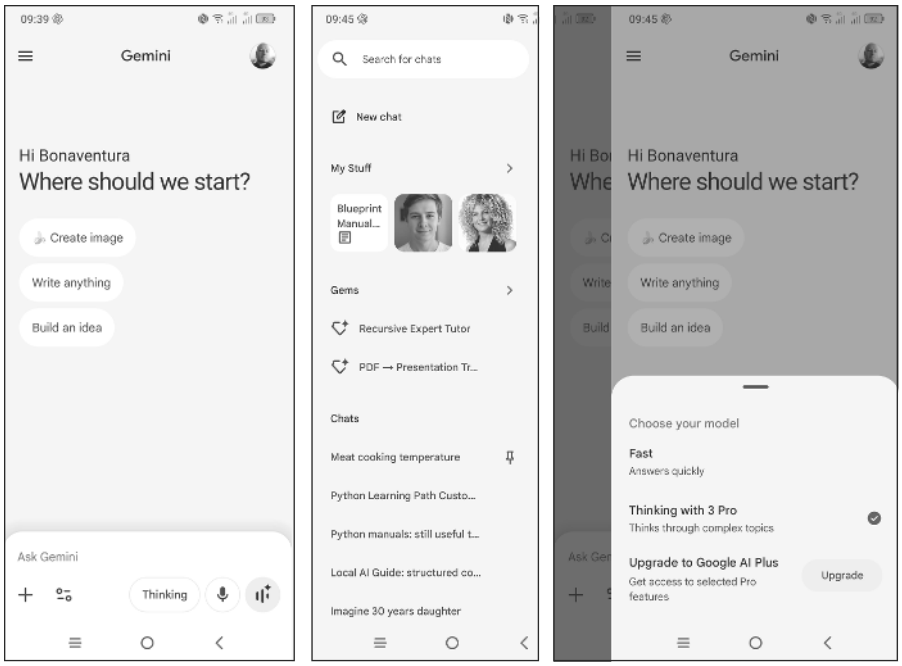


(a) (b) (c)

**FIGURE 1-27:** Three screenshots of the Gemini mobile app (Android version), including the main screen (a), with the expanded sidebar (b) and the available tools for the prompt (c).

As you can see, the mobile app is almost identical to the web (and desktop) version in its use, excluding, of course, some specific features that I discuss in more detail at the end of Chapter 4.